INTERNATIONAL JOURNAL OF ECONOMICS FINANCE & MANAGEMENT SCIENCE

Volume11 Issue01, January-2026, pg. 14-27

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

Event-Driven And Serverless Data Warehousing Architectures For Cloud-Native Analytics: Integrating Microservices, Edge Processing, And Amazon Redshift-Centric Design

Dr. Bruno Varela

Faculty of Engineering, University of Tartu, Estonia

Abstract: The contemporary data landscape is increasingly defined by unprecedented volumes, velocities, and varieties of digital traces produced by distributed applications, Internet-of-Things ecosystems, media platforms, and enterprise microservices. These conditions have exposed the limitations of traditional, monolithic data warehousing architectures that were designed for batch-oriented, centrally managed, and relatively predictable workloads. In response, cloud-native data warehousing has emerged as a paradigm that integrates elastic compute, object-based storage, and highly automated orchestration mechanisms, while increasingly relying on serverless and event-driven architectures to cope with real-time analytical demands. This article develops an extensive theoretical and empirical exploration of how event-driven and serverless computing models transform the design, operation, and governance of modern data warehouses, with particular attention to Amazon Redshift as a representative industrial platform for large-scale analytical processing (Worlikar, Patel, & Challa, 2025).

Drawing upon an interdisciplinary corpus of literature on serverless computing, microservices, edge architectures, and big data analytics, this study synthesizes insights from both foundational and recent research to articulate a unified conceptual framework for cloud-native, event-driven data warehousing. The article situates serverless functions and event brokers not merely as auxiliary components but as epistemic infrastructures that reshape how data is ingested, transformed, and consumed across distributed analytical ecosystems. By linking the theoretical insights of the serverless trilemma, which highlights trade-offs between composition, scalability, and performance (Baldini et al., 2017), with empirical findings on event-driven microservice coordination and technical debt (de Toledo et al., 2021; Laigner et al., 2024), the study demonstrates how architectural decisions propagate into long-term data governance and analytical reliability.

Methodologically, the article adopts a qualitative, literature-driven analytical approach that systematically integrates technical, organizational, and socio-economic dimensions of data warehousing. Rather than offering a narrow case study, the analysis constructs a generalized

Volume11 Issue01, January-2026, pg. 14-27

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

interpretive model that explains how platforms such as Amazon Redshift are increasingly embedded within event-driven pipelines that include serverless ingestion, edge-based preprocessing, and asynchronous microservice orchestration (Baresi et al., 2017; Kanso & Youssef, 2017). The results reveal that the combination of Redshift's columnar, massively parallel processing engine with serverless and event-driven front-ends enables a new form of analytical elasticity, in which data warehouses become continuously adaptive rather than statically provisioned (Worlikar et al., 2025).

The discussion critically evaluates the benefits and limitations of this architectural convergence. While event-driven and serverless designs offer unprecedented scalability, cost efficiency, and responsiveness, they also introduce new forms of complexity, opacity, and architectural technical debt that challenge traditional governance and optimization strategies (Chavan, 2021; de Toledo et al., 2021). The article concludes that the future of data warehousing will be neither purely serverless nor purely stateful, but rather a hybrid ecosystem in which platforms like Amazon Redshift operate as stable analytical cores surrounded by highly dynamic, event-driven peripheries. By providing a deeply elaborated theoretical and practical account, this study contributes to a more nuanced understanding of how cloud-native analytics infrastructures can be designed to support the next generation of data-intensive applications.

Keywords

Cloud-native data warehousing, Event-driven architecture, Serverless computing, Microservices analytics, Amazon Redshift, Big data integration

INTRODUCTION

The evolution of data warehousing has always been inseparable from broader transformations in information technology, organizational practice, and socio-economic structures. Early data warehouses emerged in an era when enterprises primarily sought to consolidate transactional data into centralized repositories that could support managerial reporting and decision-making. These systems were designed around batch processing, relational schemas, and tightly controlled extract-transform-load pipelines, reflecting both the technological constraints and the organizational hierarchies of their time. However, the contemporary digital economy, characterized by ubiquitous connectivity, real-time interaction, and algorithmic mediation, has profoundly altered the nature of data itself. Digital traces are no longer generated primarily by internal enterprise systems but by a heterogeneous constellation of mobile devices, social platforms, sensors, and microservice-based applications, each producing continuous streams of semi-structured and unstructured data (Al-Ali et al., 2017; Gessert et al., 2017). This shift has

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

exposed the inadequacy of traditional data warehousing models, which struggle to ingest, process, and analyze data at the scale and speed demanded by modern analytical workloads.

Cloud computing has been widely recognized as a critical enabler of the new generation of data warehouses because it decouples storage and compute, provides elastic resource allocation, and allows organizations to pay only for what they use. Platforms such as Amazon Redshift epitomize this transformation by offering massively parallel processing engines optimized for columnar storage and large-scale analytical queries, while abstracting away much of the infrastructure management that historically burdened data warehouse administrators (Worlikar et al., 2025). Yet even cloud-native data warehouses face significant challenges when confronted with highly dynamic, event-driven data sources. The rise of microservices, serverless functions, and edge computing means that data is often generated and transformed in ephemeral execution contexts that do not conform to the assumptions of stable, long-lived database connections or predictable batch windows (Baldini et al., 2017; Baresi et al., 2017). As a result, the architectural question is no longer simply how to scale a data warehouse in the cloud, but how to integrate it into a broader event-driven ecosystem that spans from edge devices to centralized analytical cores.

Event-driven architecture has emerged as a dominant paradigm for coordinating distributed systems in which components communicate through asynchronous messages or events rather than synchronous calls. In microservice-based applications, this approach is often justified on the grounds that it increases decoupling, resilience, and scalability, allowing individual services to evolve independently without breaking the entire system (Chavan, 2021; Gupta, 2018). However, when applied to data warehousing, event-driven design has deeper epistemological implications. It transforms the very notion of what constitutes a dataset, replacing static snapshots with continuously evolving streams of events that represent changes in the state of the world. This shift aligns closely with the increasing importance of real-time analytics in domains such as smart homes, video streaming, and online gaming, where decisions must be made on the basis of up-to-date information rather than historical aggregates (Al-Ali et al., 2017; Garcia et al., 2010; Lundberg, 2017).

Serverless computing further intensifies this transformation by providing a model in which developers write functions that are executed on demand in response to events, without the need to provision or manage servers. From the perspective of data warehousing, serverless functions can serve as highly scalable and cost-efficient mechanisms for ingesting, filtering, and transforming data before it is loaded into analytical stores such as Amazon Redshift (Kanso & Youssef, 2017; Lynn et al., 2017). The appeal of this approach lies in its promise of infinite scalability and fine-grained billing, which seem ideally suited to the bursty and unpredictable nature of event streams. Yet the literature on serverless computing also highlights significant challenges, including cold-start latency, limited execution duration, and difficulties in composing complex workflows, collectively described as the serverless trilemma (Baldini et al., 2017). These issues become particularly salient when serverless functions are used as integral components of data pipelines, where reliability, consistency, and performance are paramount.

Despite the growing importance of these technologies, the academic and practitioner literatures have often treated data warehousing, serverless computing, and event-driven microservices as largely separate domains of inquiry. Research on data warehousing has traditionally focused on query optimization, storage formats, and extract-transform-load processes, while research on serverless and microservices has emphasized application development, scalability, and operational concerns (Gessert et al., 2017; Lynn et al., 2017). This fragmentation has resulted in a significant literature gap: there is still no comprehensive theoretical framework that explains how event-driven and serverless architectures reshape the design and governance of cloud-native data warehouses. The present study seeks to address this gap by developing an integrated perspective that situates platforms like Amazon Redshift within a broader ecosystem of event-driven analytics (Worlikar et al., 2025).

The choice of Amazon Redshift as a focal point is motivated not only by its widespread industrial adoption but also by its architectural features, which exemplify many of the tensions and opportunities inherent in cloud-native data warehousing. Redshift combines columnar storage, massively parallel query execution, and deep integration with cloud object storage, enabling it to handle petabyte-scale datasets with relatively low administrative overhead (Worlikar et al., 2025). At the same time, Redshift is increasingly used in conjunction with serverless ingestion services, event brokers, and microservice-based data producers, creating complex data flows that challenge traditional notions of schema management, data quality, and workload optimization. Understanding how these components interact is therefore essential for both researchers and practitioners who seek to design robust and future-proof analytical infrastructures.

From a theoretical standpoint, the integration of event-driven and serverless paradigms into data warehousing raises fundamental questions about the nature of analytical knowledge. In classical data warehouse theory, data is often conceptualized as a stable, curated representation of organizational reality, carefully modeled and periodically refreshed to ensure consistency and reliability. In an event-driven world, by contrast, data is inherently provisional, reflecting a continuous stream of changes that may arrive out of order, be duplicated, or be revised after the fact (Cabane & Farias, 2024; Laigner et al., 2024). This epistemic instability challenges the assumption that a data warehouse can serve as a single source of truth, suggesting instead that it must be understood as one node in a distributed network of data-producing and data-consuming services. Platforms like Amazon Redshift thus become less like static repositories and more like dynamic analytical hubs that mediate between real-time event streams and long-term historical storage (Worlikar et al., 2025).

The introduction of edge computing further complicates this picture by relocating parts of the data processing pipeline closer to the sources of data generation. In smart home systems, for example, sensors and local gateways may perform preliminary analytics before sending aggregated or filtered data to the cloud, reducing latency and bandwidth consumption while improving privacy and resilience (Al-Ali et al., 2017; Baresi et al., 2017). When such edge-processed data eventually reaches a central data warehouse, it may already embody a set of interpretive decisions that shape how it can be analyzed and understood.

Integrating edge-based preprocessing with cloud-based warehousing therefore requires careful architectural and methodological consideration, particularly in event-driven environments where data flows are highly dynamic.

The growing reliance on microservices as the organizational unit of software development also has profound implications for data warehousing. Microservices encourage teams to own their data and expose it through well-defined interfaces, often using asynchronous events to communicate changes in state (Gupta, 2018; Chavan, 2021). While this approach can enhance agility and scalability, it also leads to a proliferation of heterogeneous data models and event schemas, which must be reconciled if a data warehouse is to provide coherent analytical views. Studies of microservice-based systems have shown that managing such heterogeneity over time can give rise to significant architectural technical debt, as ad hoc integration patterns accumulate and become increasingly difficult to maintain (de Toledo et al., 2021). In the context of a data warehouse, this debt manifests as brittle ingestion pipelines, inconsistent data definitions, and opaque transformation logic, all of which undermine analytical trustworthiness.

Against this backdrop, the present article advances a comprehensive research agenda that seeks to understand how event-driven and serverless architectures can be harnessed to enhance, rather than undermine, the capabilities of cloud-native data warehouses. By synthesizing insights from the diverse literatures represented in the provided references, the study develops a nuanced account of both the opportunities and the risks associated with this architectural convergence. In doing so, it responds directly to calls in the serverless and microservices communities for more empirical and theoretical work on the long-term implications of event-driven design (Baldini et al., 2017; Lynn et al., 2017; Laigner et al., 2024), while grounding its analysis in the concrete practices of data warehousing as exemplified by Amazon Redshift (Worlikar et al., 2025).

The remainder of this article is organized around a detailed methodological exposition, an interpretive presentation of results, and an extensive discussion that situates those results within broader scholarly debates. Throughout, every analytical claim is anchored in the existing literature, ensuring that the argument remains both theoretically rigorous and empirically informed. By maintaining this dual commitment, the study aims to contribute not only to academic understanding but also to the practical design of next-generation data warehousing systems that are capable of supporting the increasingly event-driven and serverless nature of digital life (Cabane & Farias, 2024; Worlikar et al., 2025).

METHODOLOGY

The methodological orientation of this study is grounded in an interpretive and integrative analysis of existing scholarly and practitioner-oriented literature on cloud-native data warehousing, serverless computing, event-driven architecture, and microservices. Rather than pursuing a narrowly scoped empirical case study or a purely theoretical exposition, the research design adopts what can be described as a qualitative synthesis methodology. This approach is particularly appropriate for domains such as

Volume11 Issue01, January-2026, pg. 14-27

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

cloud computing and data warehousing, where rapid technological change and the complexity of socio-technical systems make it difficult to isolate variables in a controlled experimental setting (Lynn et al., 2017; Baldini et al., 2017). By systematically examining and interrelating the concepts, findings, and debates contained in the provided references, the study constructs a comprehensive analytical framework that elucidates how event-driven and serverless paradigms reshape the architecture and governance of platforms such as Amazon Redshift (Worlikar et al., 2025).

At the core of this methodology lies a structured process of thematic coding and conceptual mapping. Each reference was examined to identify its primary contributions to understanding distributed systems, data processing, or analytical infrastructures. For example, works on serverless computing were analyzed for their treatment of scalability, function composition, and performance trade-offs (Baldini et al., 2017; Kanso & Youssef, 2017), while studies of microservices and event-driven architectures were coded for their insights into system coordination, technical debt, and organizational impact (Chavan, 2021; de Toledo et al., 2021; Laigner et al., 2024). Research on application domains such as smart homes, video delivery, and gaming was similarly examined to understand how real-time data generation and consumption impose specific requirements on data warehousing and analytics (Al-Ali et al., 2017; Garcia et al., 2010; Lundberg, 2017).

These thematic codes were then organized into higher-level conceptual categories that reflect the key dimensions of cloud-native, event-driven data warehousing. Among these dimensions are architectural elasticity, data pipeline dynamism, governance and technical debt, and epistemic stability of analytical datasets. The purpose of this categorization was not merely to summarize the literature but to reveal underlying patterns and tensions that cut across different technological and organizational contexts (Gessert et al., 2017; Cabane & Farias, 2024). For instance, the tension between the scalability promised by serverless functions and the composability challenges identified in the serverless trilemma becomes particularly salient when those functions are used to orchestrate complex data transformations before loading data into Amazon Redshift (Baldini et al., 2017; Worlikar et al., 2025).

A critical aspect of the methodology is its deliberate emphasis on triangulation across different types of sources. The reference list includes peer-reviewed journal articles, conference proceedings, technical reports, and practitioner-oriented analyses. Each of these genres offers a distinct perspective on the phenomena under study. Academic articles tend to provide rigorous theoretical frameworks and empirical evidence, while practitioner sources often capture emerging best practices and real-world challenges that have not yet been fully theorized (Gupta, 2018; Chavan, 2021). By juxtaposing these perspectives, the study seeks to avoid the limitations of relying on any single epistemic community and to produce a more holistic understanding of event-driven data warehousing (Lynn et al., 2017; Worlikar et al., 2025).

The methodological framework also draws on the notion of architectural narratives as a way of understanding how technologies are adopted and interpreted within organizations. In this view, architectures such as microservices or serverless platforms are not merely technical configurations but

Published Date: - 05-01-2026
E-ISSN: 2536-7897
P-ISSN: 2536-7889

also stories that articulate particular values, such as agility, scalability, or cost efficiency (Ian Rudd, 2009; Chavan, 2021). By analyzing how these narratives are reflected in the literature, the study can identify implicit assumptions that shape design decisions and, by extension, the structure of data warehousing systems. For example, the narrative of infinite scalability often associated with serverless computing may lead organizations to underestimate the long-term complexity of managing event-driven data pipelines, thereby contributing to architectural technical debt (de Toledo et al., 2021; Baldini et al., 2017).

Within this interpretive framework, Amazon Redshift serves as an anchoring case that grounds the analysis in a concrete technological context. The Amazon Redshift Cookbook provides detailed insights into how modern data warehousing solutions are built, optimized, and integrated with other cloud services (Worlikar et al., 2025). Rather than treating Redshift as an isolated product, the methodology situates it within a broader ecosystem of serverless ingestion services, event brokers, and microservice-based data producers. This contextualization allows the study to examine how theoretical concepts such as event-driven processing and serverless composition are instantiated in real-world analytical architectures (Gessert et al., 2017; Lynn et al., 2017).

The limitations of this methodology must also be acknowledged. Because the study relies on existing literature rather than primary empirical data, its conclusions are necessarily constrained by the scope, quality, and biases of the available sources. For example, much of the research on serverless computing and microservices is still exploratory in nature, reflecting the relative novelty of these paradigms (Baldini et al., 2017; Kanso & Youssef, 2017). Similarly, practitioner-oriented sources may emphasize success stories and best practices while underreporting failures or long-term challenges (Gupta, 2018; Chavan, 2021). To mitigate these limitations, the analysis explicitly engages with counter-arguments and critical perspectives, such as those highlighting the accumulation of technical debt and the performance trade-offs inherent in event-driven architectures (de Toledo et al., 2021; Cabane & Farias, 2024).

Another methodological constraint arises from the rapid pace of technological change in cloud computing. Architectural patterns that are considered best practice today may be superseded by new tools or paradigms in the near future. Nevertheless, by focusing on underlying principles such as decoupling, elasticity, and event-driven coordination, the study aims to produce insights that remain relevant even as specific technologies evolve (Lynn et al., 2017; Worlikar et al., 2025). In this sense, the methodology prioritizes conceptual robustness over temporal specificity, seeking to contribute to a durable theoretical understanding of cloud-native data warehousing.

Finally, the interpretive nature of the methodology means that the researcher’s own analytical judgments play a significant role in shaping the narrative. This subjectivity is not a flaw but a recognized feature of qualitative synthesis, provided that it is exercised transparently and grounded in the literature (Gessert et al., 2017; Laigner et al., 2024). By consistently citing the sources that inform each analytical claim, the study ensures that readers can trace the reasoning back to its empirical and theoretical foundations. In doing so, the methodology aligns with the broader scholarly commitment to rigor, reflexivity, and critical

Volume11 Issue01, January-2026, pg. 14-27

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

engagement that characterizes high-quality research in information systems and cloud computing (Baldini et al., 2017; Worlikar et al., 2025).

RESULTS

The interpretive synthesis of the literature reveals a set of interrelated findings that collectively illuminate how event-driven and serverless architectures reshape the design, operation, and epistemic role of cloud-native data warehouses, particularly when anchored around a platform such as Amazon Redshift (Worlikar et al., 2025). These results do not represent statistical measurements in the conventional sense but rather analytically grounded patterns that emerge consistently across diverse empirical and theoretical studies (Baldini et al., 2017; Lynn et al., 2017). In this regard, the findings can be understood as explanatory propositions about how distributed, event-driven infrastructures function when integrated with large-scale analytical systems.

A first major result concerns the transformation of data ingestion from a batch-oriented to an event-centric process. Traditional data warehouses were historically built around periodic extract-transform-load operations, which assumed that data sources could be polled at regular intervals and that their state at the time of extraction was sufficiently stable for analytical purposes. In contrast, the literature on microservices and event-driven systems shows that contemporary applications increasingly emit streams of events that reflect every change in application state, from user interactions to system-level updates (Gupta, 2018; Chavan, 2021; Laigner et al., 2024). When these event streams are connected to cloud data warehouses through serverless ingestion pipelines, the warehouse ceases to be a repository of snapshots and instead becomes a continuously updated reflection of organizational and user activity. Platforms such as Amazon Redshift, when combined with cloud-native streaming and serverless services, are therefore positioned not merely as storage engines but as real-time analytical substrates capable of supporting near-instantaneous insight generation (Worlikar et al., 2025; Gessert et al., 2017).

The literature further indicates that this shift to event-driven ingestion has profound implications for data modeling and schema management. In microservice-based systems, each service often maintains its own data model, emitting events that reflect domain-specific concepts and structures (Ian Rudd, 2009; de Toledo et al., 2021). When these heterogeneous events are ingested into a centralized warehouse such as Redshift, they must be reconciled into a coherent analytical schema. This process is rarely trivial, as it requires not only technical transformation but also organizational agreement about the meaning and ownership of data fields (Laigner et al., 2024). The result is that schema evolution becomes a continuous and negotiated process rather than a one-time design decision, reinforcing the idea that cloud-native data warehouses are socio-technical systems rather than purely technical artifacts (Worlikar et al., 2025; Lynn et al., 2017).

A second major result concerns the role of serverless computing in mediating between event streams and analytical storage. Serverless functions are frequently deployed as lightweight, on-demand processors that validate, enrich, and route incoming events before they reach the data warehouse (Kanso & Youssef, 2017; Baldini et al., 2017). The literature consistently finds that this approach offers significant advantages in terms of scalability and cost efficiency, as compute resources are allocated only when events actually arrive. In high-volume environments such as smart home systems or video delivery platforms, this elasticity is critical for handling bursts of activity without overprovisioning infrastructure (Al-Ali et al., 2017; Garcia et al., 2010). When integrated with Amazon Redshift, serverless ingestion pipelines allow organizations to maintain a relatively stable analytical core while dynamically scaling the periphery of their data processing architecture (Worlikar et al., 2025; Gessert et al., 2017).

At the same time, the results highlight that serverless architectures introduce new forms of complexity that directly affect data warehousing. The serverless trilemma, which identifies trade-offs between function composition, performance, and scalability, becomes particularly acute when functions are chained together to implement multi-step data transformations (Baldini et al., 2017). In such scenarios, latency can accumulate across function invocations, and error handling becomes distributed across multiple ephemeral execution contexts. For a data warehouse like Redshift, which is designed to support consistent and reliable analytical queries, this upstream volatility can manifest as delayed, duplicated, or incomplete data, undermining analytical accuracy (Worlikar et al., 2025; Cabane & Farias, 2024).

A third result pertains to the impact of event-driven microservices on data quality and governance. The literature on microservices emphasizes the benefits of decentralized data ownership, which allows teams to evolve their services independently and at their own pace (Gupta, 2018; Chavan, 2021). However, when these independently evolving services feed data into a shared analytical platform, their divergent update cycles and semantic interpretations can lead to inconsistencies that are difficult to detect and resolve (de Toledo et al., 2021; Laigner et al., 2024). The interpretive analysis suggests that cloud-native data warehouses must therefore incorporate governance mechanisms that are themselves event-driven, enabling them to detect and respond to schema changes, data anomalies, and integration failures in near real time (Worlikar et al., 2025; Lynn et al., 2017).

Edge computing emerges in the literature as a further layer of complexity and opportunity. By processing data closer to its source, edge architectures can reduce latency and bandwidth usage while enabling real-time responsiveness in domains such as smart homes and gaming (Al-Ali et al., 2017; Lundberg, 2017; Baresi et al., 2017). When edge-processed data is forwarded to a central data warehouse, it often arrives in a pre-aggregated or filtered form, reflecting local analytical decisions made at the edge. This means that the data warehouse must be capable of integrating data that embodies multiple layers of interpretation, challenging the traditional assumption that raw data is always available for retrospective analysis (Worlikar et al., 2025; Gessert et al., 2017). The result is a more distributed epistemology of data, in which knowledge is co-produced by edge devices, serverless functions, and centralized analytical engines.

Volume11 Issue01, January-2026, pg. 14-27

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

Finally, the results point to a growing tension between architectural agility and long-term sustainability. Event-driven and serverless architectures are widely praised for their ability to support rapid innovation and scaling, but the literature also documents how these same properties can lead to the accumulation of architectural technical debt over time (de Toledo et al., 2021; Laigner et al., 2024). In the context of data warehousing, this debt often takes the form of ad hoc ingestion pipelines, undocumented transformation logic, and proliferating event schemas, all of which make it increasingly difficult to maintain a coherent analytical environment (Worlikar et al., 2025; Cabane & Farias, 2024). The interpretive synthesis therefore suggests that the benefits of event-driven and serverless designs must be weighed against their potential to erode the long-term reliability and transparency of analytical systems.

DISCUSSION

The results of this study invite a deeper theoretical reflection on the evolving role of data warehouses in an era defined by event-driven and serverless computing. Far from being merely a technical upgrade, the integration of these paradigms represents a fundamental reconfiguration of how organizations produce, manage, and interpret data. At the center of this reconfiguration lies a tension between the promise of real-time, scalable analytics and the enduring need for stable, trustworthy, and governable data repositories. Amazon Redshift, as a paradigmatic cloud-native data warehouse, exemplifies both the opportunities and the contradictions inherent in this transformation (Worlikar et al., 2025).

From a historical perspective, data warehousing has always been shaped by prevailing computational paradigms. The rise of relational databases in the late twentieth century supported centralized, schema-driven warehouses optimized for batch reporting. The advent of distributed file systems and big data platforms in the early twenty-first century enabled more flexible and scalable forms of analytics, but often at the cost of increased operational complexity (Gessert et al., 2017). The current shift toward serverless and event-driven architectures can be seen as the latest phase in this evolution, reflecting a broader move toward fine-grained, on-demand computation and asynchronous coordination (Baldini et al., 2017; Kanso & Youssef, 2017). In this sense, platforms like Redshift are not static endpoints but dynamic nodes in an ever-expanding network of data-producing and data-consuming services.

One of the most significant theoretical implications of this shift is the redefinition of what it means for a data warehouse to be "central." In classical architectures, the warehouse was the definitive repository of organizational data, against which all analytical queries were run. In an event-driven ecosystem, however, data is continuously in motion, flowing through microservices, serverless functions, and edge devices before eventually being consolidated in a central store (Chavan, 2021; Baresi et al., 2017). The warehouse thus becomes one layer among many in a distributed analytical stack, raising questions about how authority, trust, and accountability are allocated across that stack (Worlikar et al., 2025; Lynn et al., 2017). This distributed epistemology challenges the notion of a single source of truth, suggesting instead that truth is negotiated and reconstructed as data moves through different architectural contexts.

INTERNATIONAL JOURNAL OF ECONOMICS FINANCE & MANAGEMENT SCIENCE

Volume11 Issue01, January-2026, pg. 14-27

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

The literature on the serverless trilemma provides a useful lens for understanding these challenges. By highlighting the trade-offs between scalability, composition, and performance, Baldini et al. (2017) implicitly point to the difficulties of using serverless functions as the backbone of complex data pipelines. While serverless architectures excel at handling large numbers of independent, event-triggered tasks, they struggle with workflows that require tight coordination, stateful processing, or low-latency responses. When such workflows are used to prepare data for ingestion into Amazon Redshift, the resulting pipelines can become fragile and difficult to reason about, undermining the reliability of downstream analytics (Worlikar et al., 2025; Cabane & Farias, 2024).

At the same time, it would be a mistake to view these challenges as purely technical. The literature on microservices and event-driven architecture emphasizes that architectural choices are deeply intertwined with organizational structures and development practices (Ian Rudd, 2009; Gupta, 2018). Decentralized teams, each responsible for their own services and data, are more likely to adopt event-driven communication patterns that prioritize local autonomy over global coherence. While this can accelerate innovation, it also makes it more difficult to maintain a consistent analytical view across the organization (de Toledo et al., 2021; Laigner et al., 2024). In this context, a data warehouse like Redshift becomes a site of negotiation between competing priorities, as teams seek to balance their own needs with the collective requirement for reliable, integrated data (Worlikar et al., 2025).

The integration of edge computing adds yet another layer to this negotiation. By enabling data to be processed closer to its source, edge architectures promise to reduce latency and improve responsiveness in time-sensitive applications such as smart homes and gaming (Al-Ali et al., 2017; Lundberg, 2017). However, they also introduce new forms of data heterogeneity, as different edge nodes may apply different preprocessing or filtering logic before forwarding data to the cloud (Baresi et al., 2017). For a central data warehouse, this means that the incoming data stream may already embody a set of interpretive choices that cannot be easily undone. The warehouse must therefore be designed not only to store data but also to document and contextualize the transformations it has undergone, a requirement that is often overlooked in purely performance-driven architectural discussions (Worlikar et al., 2025; Gessert et al., 2017).

The concept of architectural technical debt provides a further critical lens through which to evaluate the long-term implications of event-driven and serverless data warehousing. De Toledo et al. (2021) show that in microservice-based systems, short-term design decisions can accumulate into significant maintenance burdens over time, as integration patterns become increasingly convoluted and difficult to refactor. When applied to data pipelines, this debt manifests as a proliferation of bespoke event schemas, transformation functions, and routing rules that obscure the lineage and meaning of data (Laigner et al., 2024). Even a powerful analytical platform like Amazon Redshift cannot compensate for such upstream complexity, as its query results are only as trustworthy as the data it receives (Worlikar et al., 2025; Cabane & Farias, 2024).

Volume11 Issue01, January-2026, pg. 14-27

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

Yet it is important to recognize that the same properties that give rise to technical debt also enable remarkable forms of innovation. Event-driven and serverless architectures make it possible to rapidly prototype new data sources, analytics, and user-facing features without the need for extensive infrastructure provisioning (Kanso & Youssef, 2017; Lynn et al., 2017). In domains such as smart energy management or video content delivery, this agility can translate directly into competitive advantage and improved user experience (Al-Ali et al., 2017; Garcia et al., 2010). The challenge, therefore, is not to reject these paradigms but to develop governance and architectural practices that harness their strengths while mitigating their risks.

One promising direction suggested by the literature is the use of event-driven governance mechanisms that mirror the dynamism of the underlying architectures. Rather than relying on periodic audits or manual schema reviews, organizations can deploy automated processes that monitor event streams for anomalies, detect schema changes, and trigger alerts or remediation workflows when inconsistencies arise (Chavan, 2021; Laigner et al., 2024). When integrated with a data warehouse like Redshift, such mechanisms can help ensure that analytical datasets remain coherent even as upstream services evolve (Worlikar et al., 2025; Gessert et al., 2017). This approach aligns with the broader trend toward infrastructure as code and policy as code, in which governance is embedded directly into the operational fabric of the system.

Another important implication of the findings is the need to rethink performance optimization in an event-driven context. Traditional data warehouse tuning focuses on query execution plans, indexing strategies, and storage formats, all of which are well supported by platforms like Redshift (Worlikar et al., 2025). In an event-driven pipeline, however, end-to-end performance also depends on the latency and reliability of serverless functions, message brokers, and edge processors (Baldini et al., 2017; Baresi et al., 2017). Optimizing such a system therefore requires a holistic view that spans multiple layers of the architecture, challenging the traditional division of labor between application developers and data warehouse administrators (Lynn et al., 2017; Cabane & Farias, 2024).

The future research directions implied by this analysis are both rich and urgent. Empirical studies are needed to measure how different event-driven and serverless design patterns affect data quality, latency, and cost in real-world data warehousing scenarios (Laigner et al., 2024; de Toledo et al., 2021). Comparative analyses of different cloud platforms and architectural approaches could shed light on best practices and trade-offs, building on the detailed operational guidance provided by sources such as the Amazon Redshift Cookbook (Worlikar et al., 2025). Finally, interdisciplinary research that integrates technical, organizational, and ethical perspectives will be essential for understanding how these architectures shape not only analytical performance but also power, accountability, and transparency in data-driven societies (Al-Ali et al., 2017; Lynn et al., 2017).

CONCLUSION

Volume11 Issue01, January-2026, pg. 14-27

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

The convergence of event-driven and serverless computing with cloud-native data warehousing represents one of the most consequential transformations in contemporary information systems. Through an extensive interpretive analysis of the relevant literature, this study has shown that platforms such as Amazon Redshift are increasingly embedded within complex ecosystems of microservices, serverless functions, and edge devices that continuously generate and process data (Worlikar et al., 2025). This integration enables unprecedented levels of analytical agility, scalability, and real-time insight, but it also introduces new forms of complexity, uncertainty, and technical debt that challenge traditional approaches to data governance and optimization (Baldini et al., 2017; de Toledo et al., 2021).

By situating Redshift within an event-driven analytical paradigm, the article has argued that the data warehouse is no longer a static repository but a dynamic hub that mediates between diverse and evolving data sources. This reconceptualization has profound implications for how organizations think about data quality, schema management, and the very notion of a single source of truth (Laigner et al., 2024; Cabane & Farias, 2024). At the same time, the analysis underscores the enduring importance of stable, high-performance analytical engines, which provide the computational foundation upon which more ephemeral, serverless components can operate (Worlikar et al., 2025; Gessert et al., 2017).

Ultimately, the future of data warehousing lies not in choosing between centralized and distributed, or between stateful and serverless, but in designing hybrid architectures that leverage the strengths of each. By drawing on the insights of the diverse literatures examined in this study, researchers and practitioners alike can move toward more robust, transparent, and adaptable analytical infrastructures that are capable of supporting the increasingly event-driven nature of the digital world (Lynn et al., 2017; Chavan, 2021; Worlikar et al., 2025).

REFERENCES

1. Garcia, A., et al. A study of transcoding on cloud environments for video content delivery. CiteSeerX.
2. Baldini, I., et al. Serverless computing: Current trends and open problems. Research Advances in Cloud Computing.
3. Worlikar, S., Patel, H., & Challa, A. Amazon Redshift Cookbook: Recipes for building modern data warehousing solutions. Packt Publishing Ltd.
4. Lynn, T., et al. A preliminary review of enterprise serverless cloud computing platforms. IEEE Xplore.
5. Al-Ali, A. R., et al. A smart home energy management system using IoT and big data analytics approach. IEEE Transactions on Consumer Electronics.
6. de Toledo, S. S., et al. Identifying architectural technical debt, principal, and interest in microservices. ScienceDirect.
7. Gupta, P. Patterns for microservices – Sync vs Async. Medium.
8. Baresi, L., et al. Empowering low-latency applications through a serverless edge computing architecture. Service-Oriented and Cloud Computing.

Volume11 Issue01, January-2026, pg. 14-27

Published Date: - 05-01-2026

E-ISSN: 2536-7897
P-ISSN: 2536-7889

9. Laigner, R., et al. An empirical study on challenges of event management in microservice architectures. arXiv.

10. Kanso, A., & Youssef, A. Serverless. Proceedings of the 2nd International Workshop on Serverless Computing.

11. Cabane, H., & Farias, K. On the impact of event-driven architecture on performance: An exploratory study. ScienceDirect.

12. Gessert, F., et al. Quaestor. Proceedings of the VLDB Endowment.

13. Chavan, A. Exploring event-driven architecture in microservices – patterns, pitfalls and best practices. ResearchGate.

14. Ian Rudd. Microservices architecture using Netflix tech stack – conceptual view. ResearchGate.

15. Lundberg, M. Evaluation of a backend for computer games using a cloud service. DIVA.

16. Baldini, I., et al. The serverless trilemma: Function composition for serverless computing. Proceedings of Onward! 2017.