INTERNATIONAL JOURNAL OF ECONOMICS FINANCE & MANAGEMENT SCIENCE

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

Designing Elastic And Cost-Efficient Cloud Data Warehouses Using Serverless Technologies

Prof. Stefan Petrov

University of Valparaíso, Chile

Abstract: The accelerating convergence between serverless computing and cloud-native data warehousing has become one of the most consequential architectural transformations in contemporary information systems. Over the past decade, enterprises have moved from monolithic, on-premise data warehouses toward elastic, cloud-hosted analytical platforms that promise virtually unlimited scalability, fine-grained cost control, and rapid innovation. In parallel, the emergence of serverless computing has introduced an execution paradigm in which infrastructure management is almost entirely abstracted away from developers, enabling applications to scale transparently while charging only for actual consumption. While these two trends have often been discussed independently, their interaction is reshaping how data warehouses are designed, optimized, governed, and economically justified. This article develops a comprehensive theoretical and empirical synthesis of serverless-enabled cloud data warehousing architectures, with particular attention to how such models influence performance predictability, cost governance, operational resilience, and analytical agility.

Grounded in the existing body of scholarly work on serverless systems and cloud data platforms, and anchored by detailed architectural guidance from Worlikar, Patel, and Challa's Amazon Redshift Cookbook (Worlikar et al., 2025), the study examines how data warehousing workloads can be decomposed into event-driven, function-based services that coexist with persistent, columnar storage engines. Drawing upon performance evaluations of production serverless environments, economic analyses of consumption-based billing, and workload characterizations from large-scale cloud providers, the paper constructs an interpretive framework that explains why certain analytical tasks benefit from serverless execution while others remain better suited to provisioned clusters (Lee et al., 2018; Shahrad et al., 2020; Adzic & Chatley, 2017). By integrating these perspectives, the article demonstrates that serverless data warehousing is not a wholesale replacement of traditional architectures but a layered, hybrid model in which ephemeral compute complements long-lived data services.

Methodologically, the research adopts a qualitative, literature-driven analytical approach, synthesizing results from peer-reviewed studies, industry-grounded technical literature, and design

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

patterns documented in modern data warehouse engineering practice. This approach enables a nuanced interpretation of performance trade-offs, particularly around cold-start latency, concurrency management, and data locality, which are critical to analytical workloads that often demand both throughput and predictability (Lloyd et al., 2018; Eismann et al., 2021). The findings indicate that when orchestrated carefully, serverless components can significantly improve the elasticity and cost efficiency of extract-transform-load processes, ad hoc query bursts, and real-time data enrichment pipelines, while persistent warehouse engines such as Amazon Redshift continue to provide optimized query execution, indexing, and storage management (Worlikar et al., 2025).

The discussion further situates these results within broader debates about cloud governance, security, and multi-cloud interoperability. Serverless architectures introduce new challenges for observability, compliance, and zero-trust security models, especially when analytical workflows span multiple providers and edge environments (Hassan et al., 2021; Enhancing cloud security…, 2016). By critically engaging with these issues, the paper argues that future data warehousing strategies must evolve from purely infrastructural considerations toward holistic, policy-driven frameworks that balance agility with control. Ultimately, the article contributes a theoretically grounded and practically informed account of how serverless computing is redefining the boundaries of cloud data warehousing, offering both scholars and practitioners a roadmap for navigating this rapidly changing landscape.

Keywords: Serverless computing; Cloud data warehousing; Amazon Redshift; Analytical workloads; Cloud economics; Data governance

INTRODUCTION

The modern enterprise is increasingly defined by its ability to collect, process, and analyze vast volumes of heterogeneous data in near real time. From transactional records and customer interactions to sensor streams and social media signals, data has become the central substrate upon which strategic, operational, and even ethical decisions are made. Historically, data warehousing emerged as the architectural response to this need, providing centralized repositories optimized for analytical queries rather than transactional workloads. Over time, these repositories evolved from tightly coupled, on-premise systems into distributed, cloud-based platforms that promise elastic scalability and reduced operational overhead. In parallel, the advent of serverless computing has introduced a radically different approach to executing code in the cloud, one that abstracts away servers entirely and bills users solely for the compute resources consumed during execution. The convergence of these two paradigms has profound implications for how analytical systems are designed, operated, and governed, a transformation

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

that is increasingly evident in platforms such as Amazon Redshift and the surrounding ecosystem of cloud services (Worlikar et al., 2025).

At a theoretical level, data warehousing and serverless computing originate from distinct traditions within computer science and information systems. Data warehouses are rooted in database theory, emphasizing schema design, query optimization, and data integrity. Serverless computing, by contrast, arises from distributed systems and cloud economics, focusing on fine-grained resource allocation, event-driven execution, and operational abstraction. Yet in practice, the boundaries between these traditions are dissolving. Contemporary analytical workloads rarely involve only long-running, batch-oriented queries. Instead, they encompass bursts of exploratory analysis, streaming data ingestion, real-time transformation, and machine learning inference, all of which exhibit highly variable and unpredictable resource demands. Serverless computing is theoretically well suited to such variability because it allows compute capacity to scale up and down automatically in response to events, without requiring administrators to provision or decommission servers manually (Baldini et al., 2017; Kounev et al., 2023).

The growing popularity of serverless computing is not merely a technological phenomenon but also an economic one. Traditional cloud infrastructures require organizations to allocate virtual machines or containers in advance, often leading to overprovisioning to handle peak loads and underutilization during off-peak periods. Serverless platforms invert this model by charging only for the milliseconds of execution time and the amount of memory actually used, thereby aligning costs more closely with business value (Adzic & Chatley, 2017). For data warehousing, where workloads can fluctuate dramatically based on reporting cycles, ad hoc queries, or sudden analytical campaigns, this economic model is particularly attractive. However, it also introduces new uncertainties, as costs become more difficult to predict and optimize without sophisticated monitoring and governance mechanisms (Hassan et al., 2021).

Within this evolving landscape, Amazon Redshift has emerged as one of the most widely adopted cloud data warehousing platforms. Its architecture combines columnar storage, massively parallel processing, and tight integration with the broader Amazon Web Services ecosystem. The Amazon Redshift Cookbook by Worlikar, Patel, and Challa provides a detailed, practitioner-oriented exploration of how such a platform can be configured and optimized for modern analytical workloads, including those that increasingly rely on serverless components for data ingestion, transformation, and orchestration (Worlikar et al., 2025). By situating Redshift within a serverless-enabled architecture, the authors implicitly acknowledge that contemporary data warehouses are no longer monolithic systems but nodes within a dynamic, event-driven network of services.

Despite the rapid adoption of these technologies, significant theoretical and empirical gaps remain in our understanding of how serverless computing fundamentally reshapes data warehousing. Much of the existing literature on serverless systems focuses on microservices, web applications, or edge computing scenarios, with relatively limited attention to analytical workloads that involve large data sets, complex queries, and stringent performance requirements (Lloyd et al., 2018; Shahrad et al., 2020). Conversely,

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

traditional data warehousing research often assumes relatively stable, provisioned compute environments and does not fully account for the transient, stateless nature of serverless functions. This disconnect has led to fragmented design practices in which serverless components are bolted onto data warehouses without a coherent architectural rationale.

The problem is not merely academic. In practice, organizations that attempt to combine serverless and data warehousing technologies often encounter unexpected performance bottlenecks, cost overruns, and governance challenges. Cold-start latency can undermine the responsiveness of analytical APIs; data movement between ephemeral functions and persistent storage can erode performance gains; and the proliferation of independently deployed functions can complicate security and compliance efforts (Lee et al., 2018; Enhancing cloud security…, 2016). These issues highlight the need for a more integrated theoretical framework that explains not only how serverless computing can be applied to data warehousing but also when and why it should be used.

This article addresses that need by developing a comprehensive, research-driven analysis of serverless-enabled cloud data warehousing architectures. Drawing on the authoritative guidance of Worlikar et al. (2025) and the broader scholarly literature on serverless systems, the study articulates a set of conceptual models that capture the interplay between persistent data services and ephemeral compute. It examines how different classes of analytical workloads map onto these models, and it explores the economic, performance, and governance implications of this mapping. In doing so, the article contributes to both theory and practice by bridging the gap between abstract serverless paradigms and the concrete realities of modern data warehouse engineering.

From a historical perspective, the current moment represents the culmination of several decades of technological evolution. Early data warehouses were tightly coupled to specific hardware platforms, with scaling achieved primarily through vertical upgrades. The transition to distributed, shared-nothing architectures introduced horizontal scalability but also increased operational complexity. Cloud computing mitigated some of this complexity by outsourcing infrastructure management, yet it retained the basic model of provisioned resources. Serverless computing represents a further abstraction, one that dissolves the notion of fixed servers altogether and replaces it with a fluid pool of compute that is allocated on demand (Kounev et al., 2023). When this abstraction is applied to data warehousing, it challenges long-standing assumptions about capacity planning, workload isolation, and even the meaning of performance.

The literature reflects both enthusiasm and skepticism about this shift. Proponents argue that serverless computing democratizes access to scalable analytics, allowing even small organizations to run sophisticated data pipelines without investing in infrastructure expertise (Eismann et al., 2021). Critics counter that the loss of control inherent in serverless platforms can lead to unpredictable behavior and vendor lock-in, particularly for data-intensive applications that depend on fine-grained performance tuning (Lloyd et al., 2018). These debates underscore the importance of empirical grounding, which is provided in part by workload studies from large cloud providers that reveal how serverless systems behave under real-world conditions (Shahrad et al., 2020).

In synthesizing these perspectives, the present study advances a central argument: serverless computing and cloud data warehousing are not competing paradigms but complementary layers within a unified analytical ecosystem. Persistent data warehouses such as Amazon Redshift provide the stable, optimized foundation for storing and querying large data sets, while serverless functions supply the elastic, event-driven compute needed to ingest, transform, and operationalize that data (Worlikar et al., 2025). The challenge for architects and researchers alike is to understand how to align these layers in ways that maximize performance, minimize cost, and uphold governance requirements. The sections that follow elaborate this argument through detailed methodological analysis, interpretive results, and theoretical discussion grounded firmly in the existing literature.

METHODOLOGY

The methodological approach adopted in this research is fundamentally interpretive and analytical, reflecting the complex, socio-technical nature of serverless-enabled cloud data warehousing. Rather than relying on a single empirical dataset or experimental deployment, the study synthesizes a broad corpus of peer-reviewed research, technical documentation, and architectural case material to construct a theoretically coherent and practically relevant account of the phenomenon. This choice is justified by the fact that serverless data warehousing is not a discrete, easily isolated system but an evolving assemblage of technologies, organizational practices, and economic models that must be understood in relation to one another (Hassan et al., 2021; Kounev et al., 2023).

At the core of the methodological design lies a structured literature analysis. Foundational works on serverless computing provide the conceptual vocabulary and performance models necessary to interpret how ephemeral compute behaves under varying workloads (Baldini et al., 2017; Lloyd et al., 2018). Studies of production serverless environments and large-scale cloud provider workloads offer empirical grounding, revealing patterns of concurrency, latency, and resource utilization that are directly relevant to analytical applications (Lee et al., 2018; Shahrad et al., 2020). These sources are complemented by economic and architectural analyses that explore how serverless pricing models and design principles influence system-level decisions (Adzic & Chatley, 2017; Eismann et al., 2021).

Crucially, the methodological framework also integrates practitioner-oriented knowledge from the Amazon Redshift Cookbook, which serves as a bridge between abstract theory and concrete implementation (Worlikar et al., 2025). This text is treated not merely as a how-to manual but as an empirical artifact that encapsulates the accumulated experience of engineers building modern data warehousing solutions in a serverless-rich environment. By analyzing the architectural patterns, optimization strategies, and integration techniques described therein, the study gains insight into how theoretical principles are operationalized in real-world systems.

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

The analytical process proceeds through several iterative stages. First, the literature is coded thematically, with particular attention to concepts such as elasticity, latency, data locality, cost predictability, and governance. These themes are then mapped onto the architectural components of a typical cloud data warehousing ecosystem, including storage engines, query processors, ingestion pipelines, and orchestration layers. Serverless functions are examined in relation to each of these components, allowing the analysis to identify points of synergy and tension (Eismann et al., 2021; Worlikar et al., 2025).

Second, the study conducts a comparative analysis of workload characteristics. Drawing on empirical findings from Shahrad et al. (2020) and Lee et al. (2018), the methodology differentiates between bursty, event-driven workloads and sustained, throughput-oriented analytical queries. This distinction is critical because it underpins the argument that serverless computing is particularly well suited to certain classes of data warehousing tasks, such as extract-transform-load operations and ad hoc query processing, while others may still benefit from provisioned, long-running clusters (Lloyd et al., 2018).

Third, the methodological framework incorporates an economic lens. By synthesizing insights from Adzic and Chatley (2017) and Hassan et al. (2021), the analysis evaluates how consumption-based billing models influence architectural decisions in data warehousing. Cost is not treated as a simple numerical variable but as a socio-technical construct shaped by organizational budgeting practices, risk tolerance, and performance expectations. The Amazon Redshift Cookbook contributes to this dimension by illustrating how cost optimization techniques, such as workload management and query tuning, interact with the broader serverless ecosystem (Worlikar et al., 2025).

The methodology also explicitly acknowledges its limitations. Because the study relies on secondary sources rather than primary experimental data, it cannot capture every nuance of specific deployment contexts. However, this limitation is mitigated by the breadth and depth of the literature surveyed, which spans multiple years, platforms, and analytical perspectives. Moreover, the interpretive approach is particularly well suited to a field that is evolving rapidly, where rigid experimental designs may quickly become obsolete (Kounev et al., 2023).

Another important aspect of the methodology is its attention to security and governance. By incorporating literature on zero-trust architectures and multi-cloud security, the study situates serverless data warehousing within a broader organizational and regulatory context (Enhancing cloud security…, 2016). This ensures that the analysis does not reduce architectural decisions to purely technical or economic considerations but recognizes the institutional constraints that shape how systems are built and operated.

In sum, the methodological strategy of this research is designed to capture the multifaceted nature of serverless-enabled cloud data warehousing. By weaving together theoretical models, empirical findings, and practitioner insights, it provides a robust foundation for the descriptive and interpretive analyses presented in the subsequent sections (Worlikar et al., 2025; Hassan et al., 2021).

RESULTS

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

The synthesis of the literature and architectural analysis reveals several interrelated patterns that characterize the operation of serverless-enabled cloud data warehouses. One of the most salient findings is the emergence of a layered computational model in which persistent, provisioned data warehouse engines coexist with ephemeral, event-driven serverless functions. This model is evident in platforms such as Amazon Redshift, where the core query processing and storage infrastructure remains continuously available, while surrounding services for data ingestion, transformation, and orchestration are increasingly implemented using serverless technologies (Worlikar et al., 2025).

A first major result concerns performance elasticity. Empirical studies of serverless environments consistently demonstrate that they can scale to handle large numbers of concurrent requests with minimal human intervention, a property that is particularly valuable for data warehousing tasks that exhibit bursty behavior (Lee et al., 2018; Shahrad et al., 2020). When applied to data ingestion pipelines, for example, serverless functions can automatically scale in response to sudden surges in incoming data, ensuring that latency remains low even under peak loads. This stands in contrast to traditional, provisioned ingestion systems, which often require manual scaling and can become bottlenecks during traffic spikes (Lloyd et al., 2018).

At the same time, the results indicate that not all aspects of data warehousing benefit equally from serverless execution. Analytical queries that scan large volumes of data and perform complex aggregations tend to be more efficiently executed on provisioned, massively parallel processing engines such as Amazon Redshift, which are optimized for data locality, caching, and query planning (Worlikar et al., 2025). Serverless functions, by virtue of their stateless and ephemeral nature, incur overheads related to data transfer and initialization that can degrade performance for such workloads (Eismann et al., 2021). This leads to a hybrid performance profile in which serverless computing excels at orchestration and lightweight processing, while persistent clusters handle heavy analytical lifting.

A second significant result pertains to cost dynamics. The literature shows that serverless computing's pay-per-use model can yield substantial savings for workloads with intermittent or unpredictable demand, as organizations are no longer required to pay for idle capacity (Adzic & Chatley, 2017; Hassan et al., 2021). In the context of data warehousing, this advantage is most pronounced for auxiliary tasks such as data cleansing, validation, and enrichment, which may run sporadically in response to data arrivals or user requests. By offloading these tasks to serverless functions, organizations can reduce the size and cost of their core warehouse clusters without sacrificing functionality (Worlikar et al., 2025).

However, the results also highlight the risk of cost opacity. Because serverless charges accrue at a fine-grained level, it can be difficult for organizations to predict and control spending, especially when analytical workflows involve complex chains of function invocations (Hassan et al., 2021). Without robust monitoring and governance, the apparent economic efficiency of serverless computing can be undermined by runaway costs triggered by poorly designed queries or unexpected workload patterns

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

(Eismann et al., 2021). This finding underscores the importance of integrating cost management tools and architectural best practices into serverless-enabled data warehousing strategies (Worlikar et al., 2025).

A third result concerns operational resilience and agility. Serverless architectures inherently support rapid deployment and modification of individual components, enabling data teams to iterate quickly on ingestion logic, transformation rules, and analytical services (Baldini et al., 2017; Kounev et al., 2023). This agility is particularly valuable in data warehousing contexts where schemas and business requirements evolve frequently. By decoupling these dynamic elements from the core warehouse engine, organizations can reduce the risk and downtime associated with changes (Worlikar et al., 2025).

Yet this same decoupling introduces new challenges for observability and debugging. Because serverless functions are distributed and ephemeral, tracing the flow of data and identifying the source of errors can be more difficult than in monolithic systems (Lee et al., 2018). The results suggest that successful serverless-enabled data warehouses require sophisticated logging, tracing, and monitoring infrastructures to maintain reliability and accountability (Hassan et al., 2021).

Finally, the analysis reveals important implications for security and governance. Serverless computing's fine-grained execution model aligns well with zero-trust security principles, as each function invocation can be authenticated and authorized independently (Enhancing cloud security…, 2016). In a data warehousing context, this enables more granular control over who can access and manipulate data at each stage of the analytical pipeline. However, it also increases the complexity of policy management, as permissions must be coordinated across a potentially large number of functions and services (Worlikar et al., 2025).

Taken together, these results depict serverless-enabled cloud data warehousing as a nuanced, hybrid architecture that offers significant benefits in elasticity, cost efficiency, and agility, while also introducing new layers of complexity in performance management, governance, and security (Kounev et al., 2023; Worlikar et al., 2025).

DISCUSSION

The results of this study invite a deeper theoretical interpretation of what it means to integrate serverless computing into the fabric of cloud data warehousing. At one level, the hybrid architectures observed can be understood as a pragmatic compromise between two competing ideals: the stability and optimization of persistent data warehouse engines, and the flexibility and elasticity of ephemeral serverless compute. However, a more nuanced analysis reveals that this compromise is not merely technical but reflects broader shifts in how organizations conceptualize data, infrastructure, and value creation (Adzic & Chatley, 2017; Kounev et al., 2023).

From a theoretical standpoint, serverless-enabled data warehousing challenges the traditional dichotomy between data and computation. In classical data warehouse theory, data is stored in relatively static

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

structures, and computation is applied through queries that are executed by a well-defined engine. Serverless computing disrupts this model by enabling computation to be triggered by data events, user actions, or external systems, effectively embedding logic throughout the data pipeline (Baldini et al., 2017). When combined with platforms such as Amazon Redshift, this leads to an architecture in which the warehouse is no longer a passive repository but an active participant in a distributed, event-driven ecosystem (Worlikar et al., 2025).

This reconceptualization has important implications for performance theory. Traditional performance models in data warehousing focus on metrics such as query latency, throughput, and resource utilization within a fixed cluster. Serverless computing introduces additional variables, including cold-start delays, invocation concurrency limits, and network overhead associated with data movement between functions and storage (Lloyd et al., 2018; Lee et al., 2018). The literature suggests that these factors can be mitigated through careful architectural design, such as pre-warming functions or colocating serverless services with data stores, but they cannot be eliminated entirely (Eismann et al., 2021). As a result, performance optimization in serverless-enabled data warehouses becomes a multi-layered problem that spans both persistent and ephemeral components.

Economically, the hybrid model reflects a shift from capital-intensive to operational-expense-driven analytics. The pay-per-use pricing of serverless computing aligns well with the increasingly dynamic nature of data-driven business, where analytical workloads may surge in response to market events or strategic initiatives (Adzic & Chatley, 2017). However, this alignment also introduces volatility into budgeting and financial planning, as costs become more closely tied to user behavior and data volume (Hassan et al., 2021). The Amazon Redshift Cookbook emphasizes the importance of workload management and query optimization in controlling these costs, suggesting that economic governance is as much a technical challenge as a managerial one (Worlikar et al., 2025).

Security and governance emerge as another critical domain of theoretical tension. Zero-trust architectures advocate for the continuous verification of every access request, a principle that maps naturally onto the fine-grained execution model of serverless functions (Enhancing cloud security…, 2016). In a data warehousing context, this enables more precise control over who can perform which operations on which data. Yet the proliferation of micro-services also increases the attack surface and the administrative burden of maintaining consistent policies (Hassan et al., 2021). This duality illustrates a broader theme in cloud computing: increased flexibility often comes at the cost of increased complexity, requiring new forms of tooling and organizational practice to manage effectively (Kounev et al., 2023).

The findings also resonate with ongoing debates about vendor lock-in and interoperability. Serverless platforms are typically tightly integrated with their underlying cloud ecosystems, making it difficult to port applications across providers (Baldini et al., 2017). When data warehouses such as Amazon Redshift are embedded within these ecosystems, the coupling becomes even stronger, potentially constraining organizations' strategic options (Worlikar et al., 2025). Some scholars argue that this lock-in is a

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

reasonable trade-off for the efficiency and innovation that cloud platforms provide, while others warn that it can lead to long-term dependency and reduced bargaining power (Eismann et al., 2021).

Looking ahead, the future of serverless-enabled data warehousing is likely to be shaped by advances in edge computing, artificial intelligence, and multi-cloud orchestration. As data sources proliferate at the edge of the network, serverless functions may increasingly act as the first point of analytical processing, filtering and aggregating data before it reaches centralized warehouses (Muralidhara & Janardhan, 2016a). This trend could further blur the boundary between data collection and data analysis, reinforcing the need for integrated architectural frameworks (Kounev et al., 2023).

At the same time, the growing sophistication of cloud-native data warehouses suggests that persistent engines will continue to play a central role in large-scale analytics. Optimizations in storage formats, query planners, and workload management, as documented by Worlikar et al. (2025), provide capabilities that are difficult to replicate in purely serverless environments. The most plausible future, therefore, is one in which serverless and provisioned components coexist in an increasingly seamless and automated fashion, guided by policies that dynamically allocate workloads to the most appropriate execution context.

In reflecting on the limitations of the present study, it is important to acknowledge that the field is evolving rapidly. New serverless platforms, pricing models, and architectural patterns continue to emerge, potentially altering the balance of advantages and disadvantages described here (Kounev et al., 2023). Nonetheless, the core insight that serverless computing and cloud data warehousing are converging in a hybrid, layered architecture appears robust across the literature reviewed (Hassan et al., 2021; Worlikar et al., 2025).

CONCLUSION

This article has argued that the integration of serverless computing into cloud data warehousing represents a fundamental shift in how analytical systems are conceived, built, and governed. By synthesizing theoretical perspectives on serverless architectures with empirical studies of performance and cost, and grounding the analysis in the practical guidance of Amazon Redshift engineering, the study has shown that serverless-enabled data warehouses are neither a simple extension of traditional models nor a wholesale replacement for them (Worlikar et al., 2025; Kounev et al., 2023). Instead, they constitute a hybrid paradigm in which persistent data services and ephemeral compute functions work together to deliver elasticity, efficiency, and agility.

The implications of this paradigm are far-reaching. For researchers, it calls for new theoretical frameworks that account for the distributed, event-driven nature of modern analytics. For practitioners, it underscores the importance of architectural literacy and governance in harnessing the benefits of serverless computing without succumbing to its pitfalls (Hassan et al., 2021; Eismann et al., 2021). As organizations continue to invest in data-driven strategies, the ability to navigate this complex landscape will be a key determinant of competitive advantage.

Volume11 Issue01, January-2026, pg. 28-38

Published Date: - 10-01-2026

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

Ultimately, serverless-enabled cloud data warehousing exemplifies a broader trend in information systems: the move toward abstraction, automation, and consumption-based models that align technological capabilities more closely with business needs. By understanding and critically engaging with this trend, scholars and practitioners alike can contribute to the development of analytical platforms that are not only more powerful, but also more adaptable, transparent, and equitable.

REFERENCES

1. Hassan, H. B., Barakat, S. A., & Sarhan, Q. I. (2021). Survey on serverless computing. Journal of Cloud Computing, 10, 39. https://doi.org/10.1186/s13677-021-00253-7
2. Muralidhara, P., & Janardhan, V. (2016). Edge and cloud integration: Optimizing latency and resource allocation for IoT applications. International Journal of Engineering and Computer Science, 5(7), 17388–17406. https://doi.org/10.18535/ijecs/v5i7.33
3. Adzic, G., & Chatley, R. (2017). Serverless Computing: Economic and Architectural Impact. In Proceedings of the 2017 11th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (pp. 1–6). https://doi.org/10.1145/3106237.3117767
4. Shahrad, M., et al. (2020). Serverless in the Wild: Characterizing and Optimizing the Serverless Workload at a Large Cloud Provider. In Proceedings of the 2020 USENIX Annual Technical Conference (pp. 205–218).
5. Enhancing cloud security: Implementing zero trust architectures in multi-cloud environments. (2016). International Journal of Scientific Research and Management, 4(9), 4636–4664. https://doi.org/10.18535/ijsrm/v4i9.22
6. Worlikar, S., Patel, H., & Challa, A. (2025). Amazon Redshift Cookbook: Recipes for building modern data warehousing solutions. Packt Publishing Ltd.
7. Lloyd, W., et al. (2018). Serverless Computing: An Investigation of Factors Influencing Microservice Performance. IEEE International Conference on Cloud Engineering (IC2E), 159–169.
8. Kounev, S., et al. (2023). Serverless Computing. Communications of the ACM, 66(5), 1–10.
9. Muralidhara, P., & Janardhan, V. (2016). Serverless computing: Evaluating performance, scalability, and cost-effectiveness for modern applications. International Journal of Engineering and Computer Science, 5(8), 17810–17834. https://doi.org/10.18535/ijecs/v5i8.64
10. Baldini, I., et al. (2017). Serverless Computing: Current Trends and Open Problems. Research Advances in Cloud Computing. Springer Singapore, 1–20. https://doi.org/10.1007/978-981-10-5026-8_1
11. Lee, H., Satyam, K., & Fox, G. (2018). Evaluation of Production Serverless Computing Environments. IEEE International Conference on Cloud Computing (CLOUD), 442–450.
12. Gartner. (2015). Gartner says 6.4 billion connected 'things' will be in use in 2016, up 30 percent from 2015. Retrieved from http://www.gartner.com/newsroom/id/3165317