International Journal of Science and Engineering Applications
Volume 14-Issue 12, 48 – 55, 2025, ISSN:- 2319 - 7560
DOI: 10.7753/IJSEA1412.1010

Cloud-Native Framework for Enterprise Business Intelligence with AI-Driven Scalability Using Snowflake and Big Query

Abdul Khaleeq Mohammed
Department of Computer Science
Governors State University
IL, USA

Mohammed Kashif
Department of IT, Jamia Millia Islamia University, New Delhi, India

Siraj Farheen Ansari
Department of Information Technology, Trine University, IN, USA

Abstract— In such a context, the enterprises of the modern world go towards cloud-native analytics. This is to match ever-increasing needs towards real-time insights, scalability of data processing, and AI-driven decision support. The aim of this paper is to introduce an integrated cloud-native Business Intelligence framework with scalable mechanisms driven by artificial intelligence using Snowflake and Google Big Query. The proposed framework covers the following challenges a business intelligence environment faces: fluctuating analytical workloads, high operational costs, fragmentation in governance, and variability in performance across different platforms.
The architecture here proposed incorporates most of the desirable features discussed above: predictive auto-scaling, AI-driven anomaly detection, dynamic resource optimization for performance efficiency, which cuts down on unnecessary compute consumption. Elasticity is provided by Mult cluster warehouses from Snowflake and serverless execution by Big Query. AI models predict the intensity of workloads to enable the system to decide autonomously how to scale up or down. The system includes automated governance, security monitoring, and compliance controls to support data integrity and privacy in multi-cloud environments.

Query performance, cost-to-performance, and system responsiveness during periods of peak demand are demonstrated by experimental investigation. AI's limited capacity for generalization and dependence on a single source are additional issues; the process is made even more difficult by the addition of more clouds. This is where the development of a durable, scalable, and intelligent enterprise business intelligence ecosystem starts. In order to construct analytics for the future, this paper explores the transformative effects of combining cloud-native data warehousing with AI-driven orchestration.

Keywords— Snowflake, Big Query, AI-driven autoscaling, predictive scaling, multi-cloud analytics, cost optimization, data governance, real-time BI, enterprise data architecture, and autonomous analytics.

I. INTRODUCTION

By producing enormous amounts of data, deploying cloud apps, and requiring real-time information to support operational and strategic choices, organizations accelerate their digital revolutions. Conventional business intelligence tools were designed to function in an on-premises, mostly monolithic environment with strict capacity limitations. They cannot keep pace with today's analytics workloads in terms of velocity and agility requirements. Businesses are quickly integrating Internet of Things devices, multi-cloud architectures, SaaS platforms, and operational data streaming. It goes without saying that BI infrastructures handcrafted by humans and statically provisioned cannot scale [1]. Driven by this reality, the move to cloud-native business intelligence architectures is accelerating, with autonomous scaling, elastic computing, decoupling of storage and computation, and AI-assisted optimization at their core.

Cloud-native data warehouses like Snowflake and Google Big Query are at the forefront of this movement. Each one offers a fully managed, highly scalable analytic engine that may reduce operational overhead while providing a wide range of advanced capabilities such as auto-clustering, materialized views, near-real-time ingestion, and native machine learning integration. Snowflake provides multi-cluster warehouses for predictable concurrency control, whilst Big Query has created a serverless, slot-based execution paradigm with on-demand elasticity and low deployment [2]. These are indeed the cloud capabilities that lay the groundwork for AI-driven orchestration, allowing the company to transition from reactive tuning to predictive and autonomous optimization.

Nowadays, the majority of businesses employ threshold-based, static, or manually established scaling rules that are unable to adjust to the unpredictable and dynamic workloads that come with using BI. Dashboard refreshes, end-of-month processing, sudden analytic spikes, ETL surges, fluctuating performance, and unpredictable costs are some of the reasons for these spikes. In particular, AI, workload forecasting, anomaly detection, and reinforcement learning-based decisioning offer transformational ways to scale intelligently and optimize resources. It brings an end-to-end Cloud-Native Framework for Enterprise BI with AI-driven scalability powered by Snowflake and Big Query [3]. It provides architecture, auto-scaler design, and an operational blueprint that empowers high-performance, cost-efficient, resilient enterprise analytics in multi-cloud environments.

II. BACKGROUND AND RELATED WORK

A. *Cloud-Native BI Trends*

Cloud-native architectures have radically changed enterprise Business Intelligence. Cloud-native BI places a high premium on containerization, microservices, and declarative infrastructure using fully managed services in building scalable, modular, resilient analytical systems [4]. Most of the modern BI platforms are predominantly using data lake house patterns, which marry the flexibility of data lakes with the performance of warehouses, while domain-oriented data mesh principles are increasingly incorporated. Each of these trends reduces the operational overhead and eliminates traditional bottlenecks associated with rigid, monolithic stacks. Besides, near-infinite scalability, multi-

International Journal of Science and Engineering Applications
Volume 14-Issue 12, 48 – 55, 2025, ISSN:- 2319 - 7560
DOI: 10.7753/IJSEA1412.1010

region data sharing, continuous integration/deployment, and policy-driven governance are supported by the cloud-native ecosystems themselves.

B. Snowflake and Big Query as Modern Cloud Analytics Engines

The architectural innovations have made Snowflake and Google Big Query the leading cloud-native analytic engines of today [5]. The multi-cluster warehouse separates storage from compute in Snowflake; clusters scale independently to handle spikes in concurrency. Its auto-suspend/ auto-resume mechanisms keep idle costs low, while robust ingestion and governance are driven by features such as Snow Pipe, Streams & Tasks, and Time Travel.

On-demand inquiries or slot-based reservations are used to supply controlled computing in Big Query's completely serverless architecture. This eliminates the need for infrastructure administration and enables quick, autonomous scalability as workload demands change. These analytical skills are extended and improved for large-scale enterprise analytics because to developments in Big Query ML, spatial processing, and integrated AI/ML APIs. Their compatibility with cloud-native BI is further supported by the fact that they both support federated queries, external tables, and lake house connectors [6].

C. AI-driven autoscaling and workload optimization

Past research in cloud resource management has revealed the inadequacy of traditional threshold-based auto-scaling, which reacts slowly and often mismanages sudden workload variations. Machine learning-driven auto-scaling makes use of forecasting, cost modelling, and reinforcement learning and has resulted in significant gains in reducing SLA violations while optimizing resource utilization [7]. Works illustrate that predictive auto-scaling can anticipate demand peaks, pre-warm resources, reduce query queuing, and minimize unnecessary scaling events. This becomes ever more applicable to BI systems because their workloads are bursty and unpredictable due to periodic refreshes of dashboards, exploratory analyses by analysts, scheduled reporting, and ETL overlaps. It requires embedding AI into cloud-native BI for autonomous performance tuning, cost governance, and adaptive resource orchestration.

Table 1. Summary of Prior Work and Cloud-Native BI Capabilities

Category	Prior Limitations	Modern Cloud-Native Capabilities	Relevance to BI Framework
BI Architectures	Monolithic, rigid, hardware-bound	Microservices, lake house, data mesh, containerization	Enables modular, scalable BI deployments
Snowflake/Big Query Evolution	Manual provisioning, limited concurrency	Elastic compute, serverless scaling, intelligent caching	Supports high concurrency and low-latency dashboards
Autoscaling Approaches	Static thresholds, reactive scaling	Predictive ML models, RL-based policies	Delivers intelligent cost/performance optimization
Resource Management Research	Over/under provisioning common	Forecast-driven capacity planning	Stabilizes BI workload performance
Governance & Security	Fragmented controls	Unified metadata, RBAC, encryption, lineage	Ensures compliance for enterprise BI

III. Principles and Requirements of Design

It follows that the integration of AI in modern systems should be performed according to a set of well-structured design principles and technical, operational, and ethical requirements for application in cybersecurity, finance, data analytics, networking, or healthcare [8]. These will assure that the system is scalable, resilient, explainable, secure, and able to respond to continuous technological changes. Performance and transparency, automation and human control, and innovation and responsible governance must all be balanced in effective AI design.

A. Scalability and Flexibility

Increased data volumes, evolving algorithms, and growing operational needs should all be able to be accommodated by AI-powered systems [9]. Scalability would enable handling a diverse set of workloads-from deep learning inference to real-time analytics-without performance degradation, while interoperability and component modularity will be important in terms of model updates, replacements, and seamless interaction with cloud, edge, and hybrid infrastructures. Long-term survivability should be based on an architecture that is easily adaptable to any new use case as that arises.

B. Security-by-Design

Rather than being an afterthought, security by design guarantees that strong defence mechanisms are built into AI systems from the start. Secure data pipelines, encrypted communication, adversarial robustness, ongoing threat monitoring capabilities, access control, identity verification, and tamper-resistant model deployment are all necessary to prevent hostile exploitation [10]. These would also make it possible to integrate automated security methods, such self-healing or anomaly detection, into a security-by-design approach, guaranteeing threat resilience.

C. Transparency and Explainability

In this sense, transparency would mean that the reasoning behind the data processing that underlies the choices and actions of the model is visible [11]. Naturally, because explainability makes it possible to identify and analyse problems, it is linked to user trust and regulatory compliance. To support each prediction, systems must be built on intrinsically interpretable models or with extra layers of explanation like SHAP or LIME. Financial security, health, and any other high-stakes area should prioritize auditability by using logging techniques to monitor decisions.

D. Reliability, Robustness, and Performance Optimization

Reliable outcomes should be produced by AI systems in a variety of operating scenarios. Fault tolerance, lowest downtime, and recovery strategy are further factors that contribute to its reliability [12]. The ability to tolerate adversarial attacks, data drift, noise, and unexpected inputs is known as robustness. Performance optimization includes low latency inferences, efficient resource use, and ongoing model modification. In real-world deployments, all of this contributes to dependability.

International Journal of Science and Engineering Applications
Volume 14-Issue 12, 48 – 55, 2025, ISSN:- 2319 - 7560
DOI: 10.7753/IJSEA1412.1010

E. Ethical, Privacy, and Compliance Requirements

Aspects of ethical compliance include fairness, bias detection, and suitable data governance [13]. Secure multiparty computation, federated learning, and differential privacy all depend on this. On the regulatory adherence, observances toward GDPR, AI Act, and cybersecurity need to be at standards.

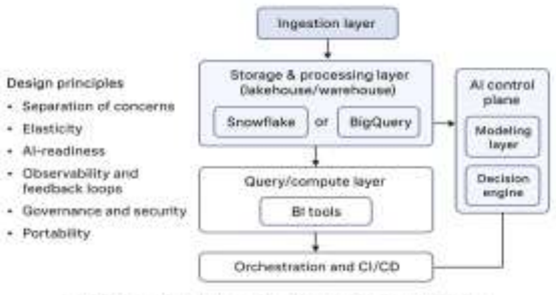

Fig. 1. Design Principles and Architecture for Cloud-Native BI

Figure 1: Design Principles Architecture for Cloud- Native BI

IV. PROPOSED FRAMEWORK: COMPONENTS AND INTERACTIONS

The Cloud-Native Enterprise BI framework's AI-driven scalability is supported by Snowflake and Big Query's modular and interoperable design [14]. Each of these will be playing a critical role in ensuring that performance optimization is carried out, real-time analytics are provided, the workload is predictable, and its execution is cost-efficient. Interactions amongst them will also be asynchronous, event-driven, and orchestrated with intelligent automation.

A. Data Ingestion & Integration Layer

This will be the layer that is charged with the tasks of batch, micro-batch, and streaming ingestion from ERP, CRM, IoT, SaaS logs, and external datasets [15]. It provides unified ingestion support, with a variety of different types of connectors-Five Tran, Kafka, Pub/Sub, and many others-to Snowflake stages or Big Query landing tables. Metadata is captured automatically to help AI-driven workload prediction downstream.

B. Cloud Data Warehouse Layer - Snowflake & Big Query

In short, Snowflake sells multi-cluster compute warehouses to scale for concurrency, while Big Query offers a serverless, distributed execution model [16]. Scaling compute and storage independent of each other is supported by both systems, enabling the AI engine to dynamically alter resource configurations. This includes performance enhancements through auto-clustering from Snowflake, among other features, and Big Query's query accelerator.

C. AI Scalability Engine

The forecasts compute the demands using time-series forecasting, anomaly detection, historical workload signatures, and reinforcement learning. It automatically scales Snowflake warehouses, such as XS → L, or allocates Big Query slots based on forecasted query intensity. Cost-conscious decision-making processes will be incorporated into the engine to attain the optimal performance/budget ratio [17].

D. Orchestration & Automation Layer

This layer is intended to control pipeline executions, track job performance, and use Airflow, Cloud Composer, or Snowflake Task to initiate AI scalability events [18]. Workflow logic now becomes vital that connects ingestion events with scaling actions and AI forecasts.

E. Layers of Compliance, Security and Governance

Both warehouses offer: role-based access, data masking, rule-based encryption, and row-level security [19]. AI detects access irregularities and automates policy.

F. The Analytics and Business Intelligence Consumption Layer

This layer includes machine learning queries, embedded analytics, and semantic models [20]. Performance feedback loops using Looker, Power BI, and Tableau by the AI Engine help with scaling optimization techniques.

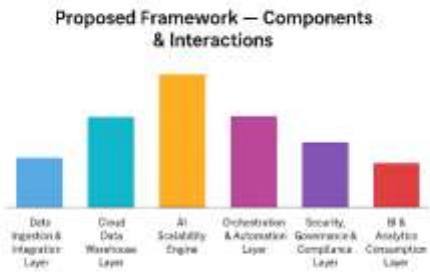

Figure 2: Proposed Framework – Components & Interactions

V. IMPLEMENTATION PATTERNS FOR SNOWFLAKE AND BIG QUERY

Adopting design patterns that put cost-effectiveness, performance, and flexibility first is required when using Snowflake or Big Query for cloud-native BI. Despite having certain basic ideas in common, they each have a different execution style. By distributing tasks intelligently, optimizing performance automatically, and allocating resources predictively, AI-driven scalability improves innate capabilities [21].

A. Patterns of Ingestion and Storage

Snowflake has links to cloud storage systems like AWS S3, Azure Blob, GCS, and many more, an ingestion pattern for continuous streaming in Snow pipe, and external stages for batch upload [22]. The internal micro-partition storage automatically optimizes the query performance itself.

Big Query supports batch ingestion via Cloud Storage, real-time ingestion via Pub/Sub, and streaming inserts with low-latency write APIs. Columnar Capacitor storage enables efficient operation over petabyte-scale datasets.

AI models account for ingestion frequency, volume of data, and bursts of events for determining when additional Snowflake warehouses or Big Query slots need to be provisioned.

B. *Compute and Query Optimization Patterns*

Snowflake offers virtual warehouses that are multi-cluster and hence can scale elastically to meet the high workloads. Some key features that enhance query performance are clustering keys, materialized views, and zero-copy cloning [23].

Big Query is powered by a serverless compute engine with automatic parallelization, dynamic slots allocation, and BI Engine acceleration [24].

AI-based auto-scaling can predict when query times are high and suggest the best warehouse sizes-for example, from Medium to Large-or update Big Query reserved slots to avoid performance bottlenecks.

C. *Cost Governance and Workload Management Patterns*

Snowflake has flexible warehouse scheduling and suspends /resume logic with per-second billing [25]. Big Query offers predictable budget control through on-demand pricing with flat-rate slot reservations.

AI continuously monitors cost anomalies, inefficient queries, and underutilized compute resources, automatically enforcing throttling rules or recommending warehouse tier adjustments.

D. *Security, Compliance, and Governance Patterns*

RBAC, dynamic data masking, encryption, and policy-based governance are supported on both platforms.

Snowflake has Time Travel and Fail-Safe, whereas Big Query natively integrates with DLP APIs for sensitive data classification.

AI-driven governance identifies abnormal access behaviours and enforces compliance policies through automated remediation [26].

Table 2: Comparison of Implementation Patterns for Snowflake vs Big Query

Category	Snowflake	Big Query
Ingestion	Snow pipe, stages, continuous loading	Pub Sub, streaming API, batch loads
Compute Model	Multi-cluster warehouses	Serverless compute, dynamic slots
Performance Features	Clustering, materialized views, micro-partitioning	BI Engine, automatic sharding, partitioning
Cost Model	Per-second billing, warehouse sizing	On-demand + flat-rate pricing
Governance	Time Travel, RBAC, zero-copy cloning	DLP API, IAM-based access
AI Scaling	Warehouse auto sizing & concurrency tuning	Slot forecasting & adaptive allocation

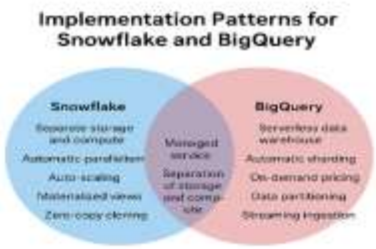

Figure 3: Implementation Patterns for Snowflake and Big Query

VI. Experimental & Evaluation Plan

This Experimental and Evaluation Plan presents efficiency, scalability, accuracy, cost-effectiveness, and reliability assessment related to a proposed AI-driven cloud-native BI framework using Snowflake and Big Query for various real-world enterprise BI workloads such as refreshes of dashboards, analytical workloads, ETL/ELT pipelines, and concurrent user activities [27]. Experiments will be executed that quantify the benefit of AI-based auto-scaling in terms of performance gains, cost reduction, and workload predictability against baseline static configurations.

A. *Experimental Setup*

These sets of experiments will be conducted in Snowflake and Google Big Query in controlled environments using identical data sets: 1TB synthetic and 500GB of real enterprise logs. Workloads include scheduled transformations, ad-hoc analytical queries, and streaming ingestion bursts.

In Snowflake, the warehouses will be set up in the XS–2XL range while Big Query will use a mix of on-demand and 2,000–20,000 slot reservations.

The AI models-LSTM, Prophet, and Reinforcement Learning agents-will be executed inside an orchestration layer that watches over workloads and decides on scaling [28].

B. *Workload Scenarios & Metrics*

a. *Performance Metrics*

- *Query latency*
- Pipeline completion time
- Concurrency overhead
- Cache hit ratio
- Autoscaling responsiveness

b. *Performance Metrics*

Calculate the total cost, including the hours and slots spent at the warehouse.

- *Cost per workload cycle*
- *Avoided over-provisioning*
- *Detecting cost anomalies with AI*

c. *Performance Metrics*

- *Failed query rate*

- *System uptime*
- *Accurate expansion at periods of high traffic.*
- *SLA compliance rate*

C. *Evaluation Procedure*

- *Phase 1: Baseline-level static scaling configurations are assessed.*
- *Threshold-based autoscaling is the second phase.*
- *The third stage could enable AI-driven predictive autoscaling.*

Each stage is repeated three times to ensure statistical validity.

Results are benchmarked separately for both Big Query and Snowflake before being included into multi-cloud systems.

Table 3: Evaluation Dimensions and Metrics

Dimension	Metric	Purpose
Performance	Latency, throughput	Assess speed under load
Cost	Compute cost, utilization	Evaluate efficiency
Reliability	Failure rate, SLA	Validate system stability
Scalability	Scaling latency, accuracy	Measure elasticity improvements
AI Quality	Prediction error, RL reward	Determine intelligence effectiveness

Figure 4: Experimental & Evaluation Plan

VII. Security, Compliance, and Governance

Strong security, regulatory compliance, and enterprise-wide governance form the cornerstones of the suggested Snowflake, Big Query, and AI-powered automation cloud-native BI solution [29]. As business intelligence systems increasingly access sensitive operational and customer data, trust, privacy, and regulatory compliance become key. This section describes how to enable enterprise-scale analytics in high-risk environments through integration of security architecture, governance practices, and compliance alignment.

A. *Security Architecture*

The solution shall be based on the Zero-Trust Security Model, which involves verifying and lawfully approving any attempt to access data and continuously monitoring it. IAM-driven rights in Big Query, VPC-Service constraints, and Customer-Managed Encryption Keys further enhance Snowflake RBAC, network limitations, and automated encryption [30].

Artificial Intelligence-powered anomaly detection techniques look at access patterns for unusual activity, credential misuse, and attempted exfiltration of data. The information is fed into real-time threat intelligence, which uses dynamic risk scores rather than conventional regulations to direct security teams' enforcement of policies.

B. *Compliance & Regulatory Controls*

Any enterprise operating across diverse geographies must adhere to a number of regulatory landscapes, such as GDPR, CCPA, HIPAA, and PCI-DSS, besides vertical-specific mandates [31]. In this respect, the proposed framework automates the compliance requirements through the embedding of data retention policies.

- Lineage and audit trails
- Access logging and immutable event records
- Automate PII/PHI classification by leveraging NLP models.

While in Snowflake, the features of governance include Snowflake Horizon; it is policy tags in Big Query that can allow automatic data classification, masking, or restriction of sensitive data within BI operations. Compliance workflows should be checked regularly by automated rule-based audits with exception reporting.

C. *Governance & Stewardship Processes*

Good governance means the consistent and reliable use of data with accountability in BI [32]. Through the use of distributed stewardship roles, a common fabric of governance is enforced by the framework. From the time data is ingested until it is visualized, data cataloguing tools monitor it. Schema, versioning, and life cycle management requirements are enforced by metadata-driven controls.

AI models can predict policy violations, recommend governance rules, and detect a decline in data quality. Automating dataset promotions, access provisioning, and schema change approvals further accelerates controlled innovation while eliminating human error.

Table 4: Security, Compliance, and Governance Mapping

Category	Key Features	Purpose
Security	RBAC, IAM, encryption, anomaly detection	Protect data from breaches and misuse
Compliance	Audits, PII masking, policy tags	Ensure adherence to regulations
Governance	Metadata, lineage, stewardship	Maintain trust, quality, and accountability

Figure 5: Layered Governance Framework

VIII. DISCUSSION: TRADE-OFFS AND PRACTICAL CONSIDERATIONS

Installing this cloud-native, AI-driven BI on Snowflake and Big Query meets basic prerequisites for a host of game-changing benefits [33]. Many typical corporate trade-offs related to cost, performance, architectural complexity, vendor dependence, skill maturity, and operational governance should be weighed against these benefits. While these practical challenges are not resolved, it is tough to come up with reasonable expectations for broad adoption.

A. *Performance vs Cost Efficiency*

AI-powered autoscaling offers great performance boosts during peak inquiries [34]. However, that very fast growth means computationally expensive processes-if proper constraints are not set in place. Big Query and Snowflake are capable of rapid expansion through their multi-cluster warehouses and on-demand autoscaling. Unexpected, ad hoc workloads may cause surges in expenses. An business needs to apply predictive workloads, set cost ceilings, and tune autoscaling thresholds in order to match financial goals with performance gains.

B. *Flexibility versus Complex Architecture*

In order to decouple computation from storage, a variety of AI services that provide dual cloud engines, flexibility, and mobility must be integrated. What that means is architectural complexity that demands mature DevOps and data engineering to keep it managed[35]. Most important of all, consistency of metadata processing across multi-cloud pipelines with unified governance policies and automated synchronization logic remains the most important factor. Less mature engineering organizations may find such complexity challenging to maintain with respect to consistent data quality, lineage, and versioning across two analytics engines.

C. *Vendor Lock-In Versus Specialized Optimizations*

Both provide highly optimized features in the micro-partitioning of Snowflake and columnar storage of Big Query with Dremel execution [36]. Using these proprietary capabilities does improve performance but at the cost of increased dependency on platform-specific constructs. A multi-cloud strategy limits lock-in risk but does require duplicate optimization and monitoring strategies. Organisations will need to consider whether the performance gain is worth the flexibility sacrificed in the longer term.

D. *Automation versus Human Governance*

AI-driven orchestration accelerates scaling decisions, anomaly detection, and resource optimization. In doing so, there is a possibility that human oversight may be reduced because of a high degree of reliance on automated decisions, particularly for environments that are sensitive to compliance. A hybrid governance model should be maintained at the enterprise level where the AI provides the intelligence, but humans retain final authority for critical access approvals, schema changes, and data quality exceptions [37].

E. *Innovation versus Operational Stability*

Innovation will be driven by predictive auto-scaling, reinforcement learning agents, and ML-based query routing, but these may introduce operational unpredictability if models misfire [38]. Rework cycles, validation pipelines, and ongoing monitoring should be used to prevent dependability issues.

IX. LIMITATIONS AND AVENUES FOR FUTURE RESEARCH

Big Query, AI-driven autoscaling, and cloud-native BI on Snowflake all improve corporate scaling, performance, and governance [39]. Interoperability, operational predictability, AI model maturation, and multi-cloud connections are some of the basic drawbacks of this relatively new architecture. As the starting point for more focused research and innovation in the future, these would be of great value.

A. *Limitations*

a. *Reliance on in-house cloud services*

Both Snowflake and Big Query have great scalability capabilities, however they employ different optimization strategies: Snowflake uses micro-partitioning, Big Query uses result caching, and Big Query uses the Dremel execution engine [40]. It is challenging to accomplish workload portability and cross-platform optimization because of these characteristics. Vendor lock-in is one of the biggest risks businesses faces when trying to offer unified multi-cloud analytics.

b. *AI Model Generalization Limitations.*

This is due to the fact that all AI models—from workload prediction to cost forecasting and anomaly detection—are based on vast volumes of historical data and have poor generalization when it comes to erratic business cycles, seasonal surges, or economic upheavals [41]. Similar to this, when faced with unknown workload distributions, reinforcement learning agents can become highly unpredictable, increasing the possibility of suboptimal scaling or unanticipated cost increases.

International Journal of Science and Engineering Applications
Volume 14-Issue 12, 48 – 55, 2025, ISSN:- 2319 - 7560
DOI: 10.7753/IJSEA1412.1010

c. Operations in multi-cloud environments are complex

Snowflake and Big Query will be able to run fluid pipelines for metadata synchronization, governance tags, a uniform lineage, and cross-cloud security controls [42]. The less developed a company's data engineering, the more difficult it is to achieve this level of operational consistency and visibility across two engines.

d. Limitations on AI Automation by Regulation

The complete autonomy of an AI-driven orchestration layer would be severely limited, thereby lowering the efficiency gains that could be achieved through automation. For instance, any scale automation or AI-driven decisions in healthcare, finance, or public services would have to be strictly supervised by humans [43].

B. Directions for Future Research

a. Cross-Cloud Semantic Optimization Layer

This semantic abstraction layer can perform the tasks of uniform rewriting of queries, optimizing execution plans, and managing lineage for future BI architecture on Snowflake and Big Query. That would greatly reduce lock-in with these vendors, and hybrid analytics would be way easier [44].

b. Federated AI Models for Workload Prediction

Federated learning allows predictive autoscaling without the need for sensitive data centralization [45]. In return, this will enhance generalization, reduce privacy risks, and enable the predictive models to learn from multi-organizational datasets.

c. Autonomous Governance Agents

It is expected that future research will be directed towards autonomous agents that can make dynamic adjustments of access privileges, detect violations of policies, and optimize rules of governance based on usage patterns [46].

d. Ethical AI and Explainability in BI Automation

Explainability modules integrated into autoscaling and anomaly detection pipelines can make decisions transparent and maintain regulatory compliance for sensitive industries [47].

Figure 6: Tradeoffs Practical Considerations

X. CONCLUSION

This paper presents the general cloud-native framework for enterprise Business Intelligence that combines AI-driven scalability with the advanced analytics features provided by Snowflake and Google Big Query. In a world where there is growing dependence on real-time insight, scalable architectures and autonomous analytics, intelligent, adaptive, and cost-efficient BI ecosystems will be increasingly critical. The proposed framework meets these requirements because it combines elastic multi-cloud data warehousing, predictive auto-scaling, automated governance, and AI-enhanced workload optimization into one single uniform architectural model.

Results from this research show that if AI-driven orchestration can predict spikes in queries, optimize resource provisioning, and reduce overprovisioning and operational inefficiencies, then dramatic gains in performance and efficiency of BI workloads are possible. Elasticity ranges seamlessly through various sets of workloads, from batch processing to near real-time analytics, using dynamic multi-cluster warehouses in Snowflake and a slot-based, scalable execution engine in Big Query. AI-driven anomaly detection, automation of governance, and predictive intelligence further help strengthen the reliability, security, and operational transparency of enterprise BI systems. It follows that limitations exist in model generalization, platform-specific optimizations, multi-cloud complexity, and regulatory constraints on automated decision-making. It thus puts into focus that balanced governance-one in which AI-driven automation supports and supplements human oversights but does not supplant them in strong areas of compliance-is necessary. Semantic multi-cloud optimization layers, federated AI models, autonomous governance agents, and explainable AI modules are some focus areas for future research to extend interoperability, trust, and decision transparency. As businesses move to fully autonomous analytics ecosystems, cloud-native BI platforms and AI-driven orchestration will define the next generation of organizational data structures, ensuring that what they produce is intelligent, scalable, and open. This demonstrates how AI-driven scaling with Snowflake and Big Query allows cloud-native BI to provide the company with a disruptive path toward high performance, adaptability, and sustained competitive advantage in data-driven decision-making.

REFERENCES

[1] Mohammed, Abdul Khaleeq, and Mohammed Azmath Ansari. "The Impact and Limitations of AI in Power BI: A."

[2] Chittoju, S. R., and Siraj Farheen Ansari. "Blockchain's Evolution in Financial Services: Enhancing Security, Transparency, and Operational Efficiency." *International Journal of Advanced Research in Computer and Communication Engineering* 13, no. 12 (2024): 1-5.

[3] Mohammed, Zubair, Naveed Uddin Mohammed Mohammed, Akheel Mohammed, Shravan Kumar Reddy Gunda, and Mohammed Azmath Ansari Ansari. "AI-Powered Energy Efficient and Sustainable Cloud Networking." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 31-36.

[4] Haven, B. "Cloud-Native Big Data Frameworks: Trends and Challenges." (2024).

[5] Thallam, Naga Surya Teja. "Comparative Analysis of Data Warehousing Solutions: AWS Redshift vs. Snowflake vs. Google BigQuery." *European Journal of Advances in Engineering and Technology* 7, no. 12 (2020): 133-141.

[6] Mohammed, Abdul Khaleeq, Siraj Farheen Ansari, Mohammed Imran Ahmed, and Zubair Ahmed Mohammed. "Boosting Decision-Making with LLM-Powered Prompts in PowerBI."

[7] Balammagary, Sruthi, Nasar Mohammed, Shanavaz Mohammed, and Asfiya Begum. "AI-Driven Behavioural Insights for Ozempic Drug Users." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 10-13.

[8] Gouni, Praveen Kumar Reddy, and Eraj Farheen Ansari. "The Impact of Cyber-Physical Attacks on AI-Enabled Business Systems."

[9] Mohammed, Naveed Uddin, Zubair Ahmed Mohammed, Shravan Kumar Reddy Gunda, Akheel Mohammed, and Moin Uddin Khaja. "Networking with AI: Optimizing Network Planning, Management, and Security through the medium of Artificial Intelligence."

[10] Mohammed, Shahnawaz, Ghousia Sultana, Fnu Mohammed Aasimuddin, and Siva Sai Ram Chittoju. "AI-Driven Automated Malware Analysis." (2025).

[11] Balasubramaniam, Nagadivya, Marjo Kauppinen, Antti Rannisto, Kari Hiekkanen, and Sari Kujala. "Transparency and explainability of AI systems: From ethical guidelines to requirements." *Information and Software Technology* 159 (2023): 107197.

[12] Keshtegar, Behrooz, Debiao Meng, Mohamed El Amine Ben Seghier, Mi Xiao, Nguyen-Thoi Trung, and Dieu Tien Bui. "A hybrid sufficient performance measure approach to improve robustness and efficiency of reliability-based design optimization." *Engineering with Computers* 37, no. 3 (2021): 1695-1708.

[13] Chittoju, Siva Sai Ram, Sireesha Kolla, Mubashir Ali Ahmed, and Abdul Raheman Mohammed. "Synergistic Integration of Blockchain and Artificial Intelligence for Robust IoT and Critical Infrastructure Security."

[14] Mohammed, Abubakar, Ghousia Sultana, Fnu Mohammed Aasimuddin, and Shahnawaz Mohammed. "Leveraging Natural Language Processing for Trade Exception Classification and Resolution in Capital Markets: A Comprehensive Study." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 14-18.

[15] Kumar, Naveen. *Intelligent Integration: Leveraging AI for Seamless ERP and CRM Connectivity*. Naveen Kumar, 2025.

[16] Nookala, Guruprasad. "Advancements in Cloud Data Warehousing: Exploring the Latest Innovations in Snowflake's Architecture and Its Impact on Data Processing Efficiency." *International Journal of Emerging Trends in Computer Science and Information Technology* 6, no. 1 (2025): 124-133.

[17] Ansari, Meraj Farheen. "Redefining Cybersecurity: Strategic Integration of Artificial Intelligence for Proactive Threat Defense and Ethical Resilience."

[18] Nookala, Guruprasad. "Snowflake’s Role in Multi-Cloud Environments: Exploring the Integration and Interoperability of Snowflake across Different Cloud Platforms." *International Journal of AI, BigData, Computational and Management Studies* 6, no. 2 (2025): 48-56.

[19] Mohammed, Akheel, Naveed Uddin Mohammed, Shravan Kumar Reddy Gunda, and Zubair Mohammed. "Fundamental Principles of Network Security."

[20] Khadri, Syed Waheeduddin, Irfan Khan Mohammed, Haroon Rasheed, and Shravan Kumar Reddy Gunda. "Adaptive Trade Exception Handling in Financial Institutions: A Reinforcement Learning Approach with Dynamic Policy Optimization." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 19-23.

[21] Ahmed, Mohammed Imran, Abdul Raheman Mohammed, Srujan Kumar Ganta, Sireesha Kolla Kolla, and Mohammed Kashif Kashif. "AI-Driven Green Construction: Optimizing Energy Efficiency, Waste Management and Security for Sustainable Buildings." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 37-41.

[22] Antoine, Lefèvre, and Garnier Sophie. "DESIGNING RESILIENT CLOUD STORAGE ARCHITECTURES: A DEEP DIVE INTO AMAZON S3, AZURE BLOB STORAGE, AND GOOGLE CLOUD STORAGE." *Journal of Engineering, Mechanics and Modern Architecture* 3, no. 9 (2024): 59-77.

[23] Soni, Ruchi. "Cloning." In *Snowflake SnowPro™ Advanced Architect Certification Companion: Hands-on Preparation and Practice*, pp. 105-112. Berkeley, CA: Apress, 2023.

[24] Janamolla, Kavitha, Ghousia Sultana Sultana, Fnu Mohammed Aasimuddin, Abdul Faisal Mohammed, and Fnu Shaik Aqheel Pasha Pasha. "Integrating Blockchain and AI for Efficient Trade Exception Handling: A Case Study in Cross-Border Settlements." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 24-30.

[25] Nookala, Guruprasad. "Advancements in Cloud Data Warehousing: Exploring the Latest Innovations in Snowflake's Architecture and Its Impact on Data Processing Efficiency." *International Journal of Emerging Trends in Computer Science and Information Technology* 6, no. 1 (2025): 124-133.

[26] Khadri, Waheeduddin, Janamolla Kavitha Reddy, Abubakar Mohammed, and T. Kiruthiga. "The Smart Banking Automation for High Rated Financial Transactions using Deep Learning." In *2024 IEEE 3rd World Conference on Applied Intelligence and Computing (AIC)*, pp. 686-692. IEEE, 2024.

[27] Faruk, Omar Muhammad. "Advanced Computing Applications in BI Dashboards: Improving Real-Time Decision Support For Global Enterprises." *International Journal of Business and Economics Insights* 4, no. 3 (2024): 25-60.

[28] Syed, Waheeduddin Khadri, Abubakar Mohammed, Janamolla Kavitha Reddy, and S. Dhanasekaran. "Biometric authentication systems in banking: A technical evaluation of security measures." In *2024 IEEE 3rd World Conference on Applied Intelligence and Computing (AIC)*, pp. 1331-1336. IEEE, 2024.

[29] Mehmood, Khawaja Tahir, Zia Ashraf, Raza Iqbal, Adnan Ahmed Rafique, Hassam Gul, and Muhammad Ali. "Cyber security Governance as a Pillar of Enterprise Risk Management: Designing a Compliance-Driven Framework for Operational Resilience, Policy Enforcement, and Regulatory Alignment." *Annual Methodological Archive Research Review* 3, no. 5 (2025): 59-77.

[30] Lopez, Gerry. "Ensuring Data Integrity and Access Control in Snowflake via RBAC and MFA."

[31] Okusi, Oluwatobiloba, Iheanacho J. Obiakor, and Fadekemi Chidinma Adeloye. "Designing resilient data risk management protocols for regulatory compliance and cyber incident response." *International Journal of Advance Research Publication and Reviews* 2, no. 7 (2025): 45-70.

[32] Pomeranz, Emily F., and Richard C. Stedman. "Measuring good governance: piloting an instrument for evaluating good governance principles." *Journal of Environmental Policy & Planning* 22, no. 3 (2020): 428-440.

[33] Kumar, Yulia, Jose Marchena, Ardalan H. Awlla, J. Jenny Li, and Hemn Barzan Abdalla. "The AI-powered evolution of big data." *Applied Sciences* 14, no. 22 (2024): 10176.

[34] Yenugula, Manideep. "Monitoring performance computing environments and autoscaling using AI." *International Journal of Computing Programming and Database Management* 4, no. 1 (2023): 90-96.

[35] Mohammed, Akheel, Zubair Ahmed Mohammed, Naveed Uddin Mohammed, Shravan Kumar Gunda, Mohammed Azmath Ansari, and Mohd Abdul Raheem. "AI-NATIVE WIRELESS NETWORKS: TRANSFORMING CONNECTIVITY, EFFICIENCY, AND AUTONOMY FOR 5G/6G AND BEYOND."." (2021).

[36] Seenivasan, Dhamotharan. "Optimizing cloud data warehousing: a deep dive into snowflake's architecture and performance." *International Journal of Advanced Research in Engineering and Technology (IJARET)* 12, no. 3 (2021): 951-962.

[37] Mohammed, Nasar, Abdul Faisal Mohammed, and Sruthi Balammagary. "Ransomware in Healthcare: Reducing Threats to Patient Care." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 2 (2025): 27-33.

[38] Awad, Mohamad Khattar, Marwa Hassan Hafez Ahmed, Ali F. Almutairi, and Imtiaz Ahmad. "Machine learning-based multipath routing for software defined networks." *Journal of Network and Systems Management* 29, no. 2 (2021): 18.

[39] Ayyub, Saim. "AI/ML-Powered Business Intelligence: Securing ERP Cloud Solutions with Snowflake DB and Advanced Cloud Computing Strategies." (2021).

[40] Levandoski, Justin, Garrett Casto, Mingge Deng, Rushabh Desai, Pavan Edara, Thibaud Hottelier, Amir Hormati et al. "BigLake: BigQuery's evolution toward a multi-cloud lakehouse." In *Companion of the 2024 International Conference on Management of Data*, pp. 334-346. 2024.

[41] Kashif, Mohammed, Mohammed Aasimuddin, Mubashir Ali Ahmed, Laxmi Bhavani Cheekatimalla, Eraj Farheen Ansari, and Ahwan Mishra. "AI-DRIVEN CTI FOR BUSINESS: EMERGING THREATS, ATTACK STRATEGIES, AND DEFENSIVE MEASURES."

[42] Dolhopolov, Anton, Arnaud Castelltort, and Anne Laurent. "Implementing federated governance in data mesh architecture." *Future Internet* 16, no. 4 (2024): 115.

[43] Mohammed, Shanavaz, Nasar Mohammed, Sruthi Balammagary, Sireesha Kolla, Srujan Kumar Ganta, and Shuaib Abdul Khader. "HARNESSING ARTIFICIAL INTELLIGENCE FOR PUBLIC HEALTH AND EPIDEMIOLOGY: OPPORTUNITIES, BARRIERS, AND PATHWAYS TO EQUITABLE GLOBAL IMPACT."

[44] Temitope, Abass. "Integrating Hybrid Vendor Management Systems with Predictive Analytics for Proactive Inventory Optimization and Lead Time Reduction in High-Pressure Supply Chains." (2025).

[45] Ahamed, Zaakki, Maher Khemakhem, Fathy Eassa, Fawaz Alsolami, Abdullah Basuhail, and Kamal Jambi. "Deep reinforcement learning for workload prediction in federated cloud environments." *Sensors* 23, no. 15 (2023): 6911.

[46] Kampik, Timotheus, Adnane Mansour, Olivier Boissier, Sabrina Kirrane, Julian Padget, Terry R. Payne, Munindar P. Singh, Valentina Tamma, and Antoine Zimmermann. "Governance of autonomous agents on the web: Challenges and opportunities." *ACM Transactions on Internet Technology* 22, no. 4 (2022): 1-31.

[47] Abi, Roland. "Ethical and Explainable AI in Data Science for Transparent Decision-Making Across Critical Business Operations."