

Scalable and Automated Business Intelligence Reporting Using Power Bi and Aws Redshift

Peeyush Patel

Oklahoma State University, Stillwater, OK, USA

Abstract: The increasing need for data-driven, real-time decision-making has necessitated the need to seek scalable and automated Business Intelligence (BI) reporting systems. The traditional systems and solutions of BI are typically not applicable to process large volumes of data, automate business processes, and deliver timely information. The combination of Microsoft Power BI and Amazon Web Services (AWS) Redshift will be discussed in this paper as the solution to this limitation. Power BI provides a powerful visualization, reporting, and automation, and AWS Redshift provides a highly scalable, cloud-based data warehouse optimized to answer complex analytics queries. With the integration of these platforms, organizations are able to make data integration processes easier, data refresh simpler, and offer interactive dashboards at a large scale. Recent literature, architecture, and other relevant issues of scalability and automation are summarized in this review. The findings indicate the value of integrating Power BI and AWS Redshift to transform raw data into usable data through effective, scalable, and automated BI reporting.

Keywords: Business Intelligence (BI), Power BI, AWS Redshift, Scalable Data Analytics, Automated Reporting, Cloud-Based BI

INTRODUCTION

BI is a mandatory part of any contemporary company. Businesses have been empowered to convert raw data into insights that are used to inform strategy decisions. In the digital age, a range of information has been created: structured and unstructured, within the context of online transactions, social media conferencing, or through IoT devices and enterprise applications (Kimball, R., & Ross, M. 2013). The ancient BI systems have been implicated by low response time of obtaining valid and accurate insight before the scaling bottlenecks, high operational costs, and their inefficiencies associated with human origins overshadow their application (Llave, M. R. 2017). In this regard, organizations are shifting to BI systems that are real-time responsive, automated, and scalable. Such systems reduce manual intervention and enhance the accuracy of the data, speed of reporting, and operational elasticity. This growing need has spawned cloud-based BI, which not only has elastic capacity but can also do on-demand storage and can readily be extended to heterogeneous information sources. Cloud BI solutions apply built-in advanced analytics, artificial intelligence, and machine learning to deliver foretelling and prescriptive insight to enable organizations to predict market trends, automate operations, and establish a culture of data-led choice throughout the company.

BI on cloud has revolutionized the world of analytics compared to the traditional constraints of locally-based management software. With affordability, scale, and access, a system like this

allows enterprises to respond rapidly to changing loads of data and provide insight into the business through web and mobile-friendly applications, allowing decision-makers to access the appropriate business data regardless of location and time. One of the leading elements or components of this architecture would be a cloud data warehouse that forms the pinnacle of massive analytics. Among the most popular proposed solutions in the Web services of cloud computing, there is Amazon Web Services, known as Redshift, which is promoted as a very high-speed analytical engine and is developed on the basis of MPP and columnar storage (Gupta, A. 2015). Redshift is petabyte-scalable, has the ability to serve petabytes of data, and is also dynamically scalable, which implies that a business may experience a rather dynamic workload (Biswas, D. 2020). The visualizing strength of the time had been augmented with the availability of easy interfacing, wide connection to the database, and robust capabilities in Microsoft Power BI, such as analysis through DAX, AI, and automatic reporting (Martins, A. *et al.*, 2020). Interactive dashboards that are based on scalability provided by Power BI and AWS Redshift provide a platform of almost real-time reporting, meeting the requirements of ecosystems, thus smoothing analytics pipelines and tuning the decision-making path (Chaalal, H. *et al.*, 2020).

Despite all the benefits of the AWS Redshift and Power BI combination, a range of problems still exist to ensure the adoption of BI ecosystems that are fully automated and scaled. The shortcomings

of query optimization in Redshift (a technical limitation), the delay in creating real-time connectivity, and the compromises between Direct Query and the Import method of interaction with Power BI are all realities that can negatively impact the work of the system. Moreover, any non-visible cost associated with repeated querying, massive block transfer, and automated recalibration will only act to degrade system efficiency, unless properly controlled, as an organization grows. Such issues also point to future research and development opportunities to advance BI models and optimize the data pipeline, and come up with other high-performance cost-balancing strategies. However, the synergies of Power BI/ AWS Redshift are a reminder of the vast opportunities to develop scalable and automated BI systems that streamline quick decision-making.

RELATED WORK

BI has gained a lot of momentum and has been under research in academia and in the industry. The earlier concepts, like dimensional modeling by Kimball, were the beginning of BI architecture research, and they unveiled scientific data warehousing and reporting methods. Later studies have examined the motivation of organizations, technology, and analysis to embrace BI or

technology to enhance decision-making, efficiency, and competitiveness. According to more recent research, this has changed to cloud-native data warehouses, columnar storage techniques, and simpler visualization technology, all of which are evidence of the BI model to emphasize scalable, automated tasking and to enable non-technical users. These developments in channels allow organizations to handle complex, large data and democratize the process of analytics. The industry developments also affect the form of the BI model.

There are also vendor-specific documents, such as white papers by Amazon Web Services (AWS) and Microsoft, which provide explicit guidance or support platform-based features, such as scalability on demand, controlled infrastructure, advanced visualization, and AI-driven data, on their platforms, Redshift and Power BI. Similarly, the scholarly work on survey literature exhibits a concentration on structures of governance, issues of implementation, and determinants of success. In summary of what is already available knowledge, Table 1 summarizes the main work of (Kimball, R., & Ross, M. 2013; Zardbani, F. *et al.*, 2020). sources under study, the focus of interests, tools, the contributions, and limitations of them, which is basically the systematic reflection of BI history.

Table 1: Summary of Literature on BI, Cloud Warehousing, and Automation

Ref	Focus Area	Platform/Tools	Key Contribution	Limitations
(Kimball, R., & Ross, M. 2013)	Dimensional modeling for BI	Data Warehouses	Defined core BI design practices for scalability	On-premise focus, less relevant for cloud BI
(Llave, M. R. 2017)	BI adoption & impact	General BI frameworks	Linked BI & analytics to organizational value	Limited focus on automation & cloud
(Gupta, A. 2015)	Cloud-native BI scaling	Amazon Redshift	Showed Redshift's MPP & scalability features	Vendor-specific, lacks broader comparison
(Biswas, D. 2020)	OLAP & DW tech	OLAP systems	Early foundation of BI querying & analysis	Outdated for modern cloud BI
(Martins, A. <i>et al.</i> , 2020)	Automated BI reporting	Power BI	Documentation of automation & refresh features	Non-academic, descriptive
(Chaalal, H. <i>et al.</i> , 2020)	Modern DW trends	Cloud DW & BI	Discussed BI challenges in scalability & governance	No specific Power BI/Redshift focus
(Gaardboe, R., & Jonasen, T. S. 2018)	Storage model optimization	Column-store DBs	Showed the efficiency of columnar storage (basis of Redshift)	Not BI-specific, no visualization focus
(Cappelli, P., & Tavis, A. 2016)	Big Data & HPC	HPC + Big Data	Highlighted convergence of big data & cloud	Not specific to BI/reporting
(Ogiela, L. <i>et al.</i> ,	BI success	Enterprise BI	Linked BI capabilities	Organizational

2020)	factors		to decision-making	view, little tech focus
(Pherson, R. H., & Heuer Jr, R. J. 2019)	Performance management	BI scorecards	Emphasized the role of BI in performance measurement	Dated, before cloud BI
(Harris, J. G., & Craig, E. 2011)	Cloud data stores	NoSQL/NewSQL	Reviewed data management in the cloud	Limited focus on BI visualization
(Armenatzoglou, N. et al., 2022)	Analytical leadership	Self-service BI	Stressed user-driven analytics adoption	Practitioner-focused, not deeply technical
(Zardbani, F. et al., 2020)	Automated BI survey	DW & Self-service BI	Surveyed the evolution from DW to automated BI	Broad survey, lacks platform-specific insights

TECHNOLOGICAL BACKGROUND

Business intelligence (BI) systems are also changing the way they are organized fundamentally due to the increased usage of cloud computing and state-of-the-art visualization systems. BI systems used in the past used hardware infrastructures and batch reporting. BI systems today are featured by distributed data storage, scalable computing resources, and reporting applications, which can be utilized with minimal training and with which it is feasible to generate insight with minimal training and in real time. These developments not only enhance scalability, efficiencies but also avail of

unhindered access to reliable data to the decision-makers.

To illustrate an example of a BI architecture, as in Fig. 1, Amazon Redshift is employed as the data warehouse in the cloud, and Microsoft Power BI is employed as the visualization and reporting layer. With this architecture, the raw data that contains incoming data in various sources, like transaction systems, web applications, and IoTs, is ingested into Redshift using Extract-Transform-Load (ETL) or Extract-Load-Transform (ELT) pipelines. Redshift optimizes the data and makes it available to sophisticated analytical queries on a large scale. Then, Power BI is directly linked to Redshift with an interactive dashboard, regularly scheduled reports, and automatic notifications.

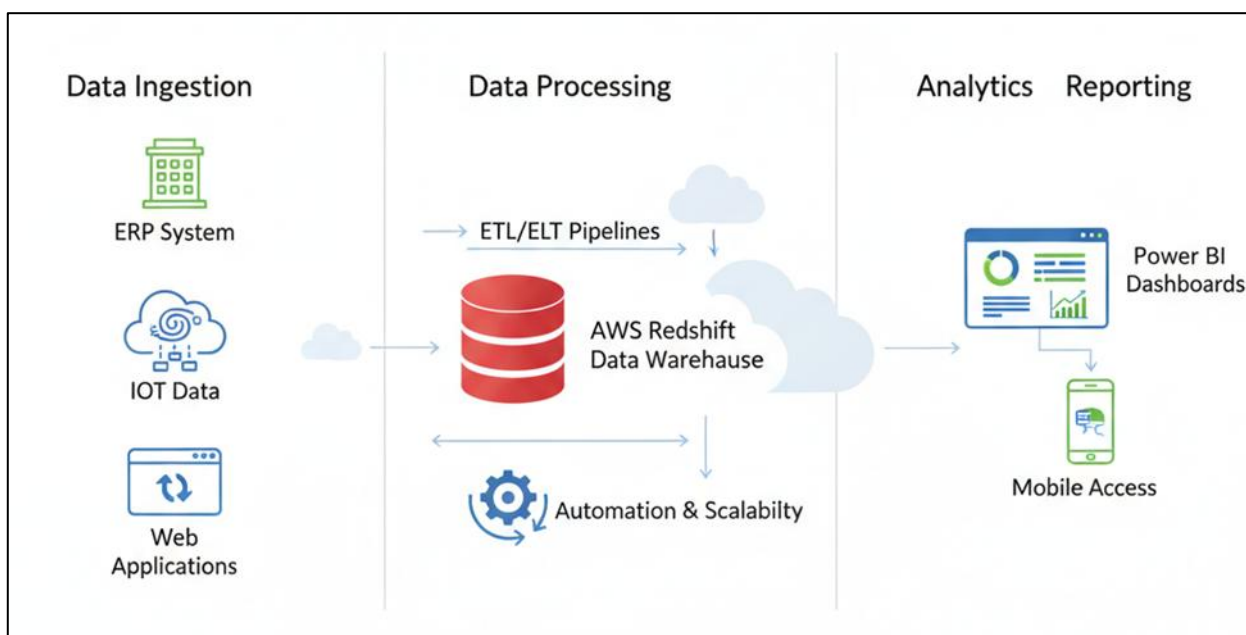


Figure 1: Conceptual architecture of scalable BI reporting

AWS Redshift

Amazon Redshift is a managed, scale-out data warehouse that is cloud-operated and optimized to

provide large-scale analytics. It is based on the architecture of Massively Parallel Processing (MPP) and columnar storage integrated to enable

the platform to effectively process the complicated analytic queries of scale datasets (petabytes scale) in an efficient manner (Adelusi, B. S. *et al.*, 2022). The MPP model distributes the data to different compute nodes, which provide a tolerable amount of concurrency and parallelism in their processing, and columnar storage reduces the input/output processing and the access to the needed attributes only (Agrawal, D. *et al.*, 2022). Its strengths include that it provides high query performance, high elastic scalability, and easy integration with other AWS services such as S3, Glue, and Kinesis that can be used to provide a variety of data ingestion and transformation needs (Pandis, I. 2021; Pavlo, A. *et al.*, 2017). However, Redshift is not perfect, particularly in cost management, since the volume of work will be proportional to the number of times the query is made by the visualization tools, which will not only change the number of computations but also raise the cost accordingly (Shadbolt, N. *et al.*, 2019). Besides these, query optimization and schema design are also significant attributes that must be utilized to ensure that performance is ensured because improperly-designed workloads can affect performance (Francis, R. *et al.*, 2023).

Microsoft Power BI

Among the self-service BI, having become popular and being based on unifying data preparation, data modeling, visualization, and AI-dependent insights into one ecosystem, is the Microsoft Power BI. It helps to perform more intensive ETL operations using Power Query, more complex computations using data analysis expressions (DAX), and natural speech querying and predictive modeling features (Bojičić, I. *et al.*, 2016) Two big methods to connect to the cloud data warehouses of Power BI, like Redshift, are Import and DirectQuery modes. The Import Mode loads the Information into the Power BI model because it is simpler to work with, but must be regularly updated since DirectQuery offers real-time information at the cost of extra latency and higher query costs to Redshift (Saxena, M. *et al.*, 2023). To make the automation possible, Power BI facilitates the scheduled refresh, coordination of power flows, and REST APIs in Power Automate to provide developers with the ability to programmatically act on datasets, triggers, and reporting lifecycles (Kulkarni, A. 2023). All that enables the companies to get out of the fixed reportage era, into the automated and dynamic BI systems.

Redshift to Power BI

Power BI allows connecting with an AWS Redshift database using native connectors and also by using ODBC/JDBC connectors, which can also be used to enable Import Mode and DirectQuery connections (Vajpayee, A. *et al.*, 2024). Connection methodology is grounded on the actual requirements of the organization in terms of real-time analysis requirements, rather than cost reduction. The other element of integration that is applicable is the data refresh under Import Mode. Power BI relies on the refresh schedules that must be duly determined to ensure timely refresh, but in the DirectQuery mode, each visualization request generates a query that must be duly configured in Redshift about the schema and indexing plans (Ahmadi, S. *et al.*, 2023). Most companies build pipelines of ETL/ELT on the AWS Glue or Lambda platform to pre-process data and inject it into the Redshift so that it can be used in the Power BI dashboards to gain better reliability and performance. This will ensure that automated reporting pipelines are scaled to ensure that they can deliver high-quality data insights to users at any time.

COMPARATIVE REVIEW OF BI ARCHITECTURES

Business Intelligence (BI) architecture has evolved from the existing on-premise systems to the new cloud-based and hybrid-based systems as a result of the growing demand for scalability, real-time information, and automation. The traditional off-premise BI solutions are defined by centralized data warehousing and the ETL procedures and reporting applications that require a significant amount of manual data entry and maintenance. On the other hand, cloud-based frameworks exploit distributed storage, elastic computing, and self-service BI systems, which enable companies to seamlessly scale resources and reduce overhead operations. Hybrid solutions refer to a combination of the two approaches, and they allow enterprises to store sensitive information in their facilities and to use cloud services to perform analytics and visualization. The relative comparison of those approaches illuminates the variability concerning the aspects of scalability, automation, cost, and performance that are salient concerns in the modern business world in demand of efficient and reliable BI solutions.

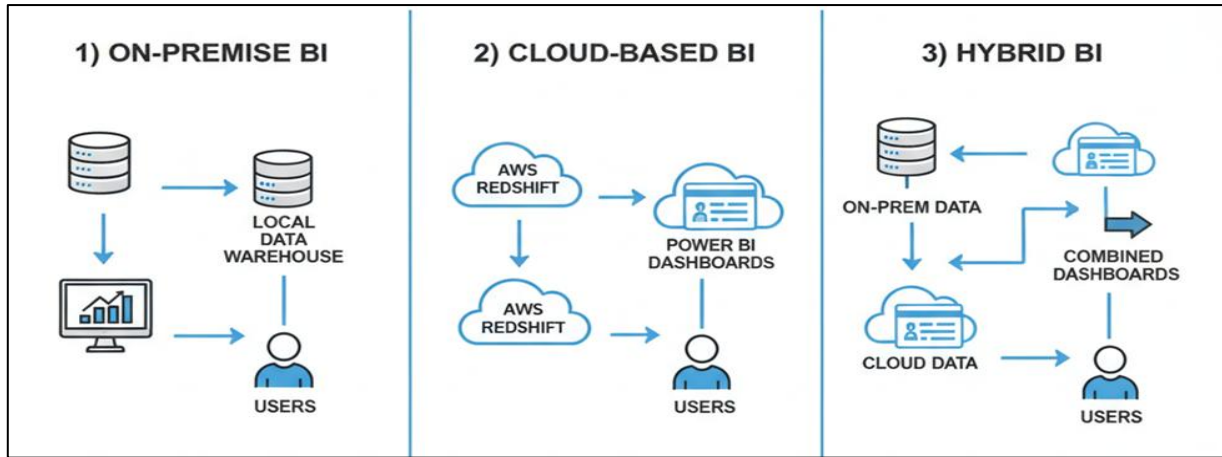


Figure 2: Comparative overview of BI architectures

Modern cloud Business Intelligence (BI) architectures, like AWS Redshift and Microsoft Power BI, are scalable, flexible, and automated, and outperform on-premise systems. These cloud-native solutions allow near-real-time processing, easy connectivity to various data sources such as IoT and web applications, automatic updates, and dynamism of computation and storage to efficiently serve fluctuating workloads. Although sensitive data can be kept on-premise with hybrid BI models, the models increase complexity in data alignment, data governance, and data latency, and need meticulous ETL planning and scheduling to prevent bottlenecks and discrepancies in reporting. The comparison of cloud, hybrid, and on-premise BI in terms of trade-offs can help organizations pick a suitable model that fits their budget, performance, scalability, and operational efficiency. BI pipelines based on cloud-first are also highly self-managed, so they require fewer staff, accelerate reporting, and enhance enterprise-wide decision-making. On-premise BI is the choice where an organization has stringent regulatory or data governance demands, like in healthcare, government, or finance. Cloud BI,

conversely, should be used by expanding businesses when they require real-time dynamic analytics without investing heavily in infrastructure, and hybrid BI should be used by larger organizations when they require the benefits of local data security and cloud-based collaboration, and scalability to maximize performance and strategic value.

CASE STUDY

Implementation of Bi Architectures in Real-World Enterprises

In order to supplement the theoretical discussion of BI architectures, we analyze practical BI systems of cloud, hybrid, and on-premise implementations. These case studies identify the achievements, advantages, and difficulties encountered by organizations and give empirical results to help make architectural decisions. Cloud BI is scalable and automated, hybrid models are more secure and flexible in analytics, and on-premise BI helps stick to stringent data regulation rules. The most important findings in these implementations are summarized in Table 2.

Table 2: Comparative Case Study of BI Architecture Implementations

BI Architecture	Example Organization	Key Outcomes	Benefits	Challenges	Cost / Observed Benefits
Cloud BI	Mid-sized E-commerce (e.g., Zalando)	5x data handling, daily reporting	Scalability, automated ETL, faster decision-making	Initial migration complexity	Reduced infrastructure cost by 40% compared to on-premise, primarily due to lower server maintenance and storage costs; faster ROI due to accelerated reporting
Hybrid BI	Multinational Bank (e.g., HSBC)	Secure sensitive data, real-	Combines the security of on-premise with the	Data integration complexity,	Moderate cost; optimized resource allocation by keeping

		time market trends	analytics flexibility of cloud; supports strategic planning	ETL scheduling	sensitive data on-premise while offloading analytics to the cloud
On-Premise BI	Regional Healthcare (e.g., Mayo Clinic regional branches)	Stable, compliant, secure	Ensures regulatory compliance, strong data privacy	Limited scalability, manual reporting	High maintenance cost due to server upkeep and IT staff; limited flexibility, but ensured compliance with healthcare regulations

SCALABILITY AND AUTOMATION CONSIDERATIONS

The scalability and automation features of the modern Business Intelligence (BI) are also rated higher because they directly influence the performance, economy, and decision-making speed. Scalability is the capability of a BI system to process larger amounts of data, simultaneous users, and complicated queries without performance deterioration. Horizontal and vertical

scaling of cloud-based data warehouses like AWS Redshift is enabled by elastic cluster resizing, concurrency scaling, and auto-scale storage, which enables businesses to effectively handle fluctuating workloads. Scalability is also concerned with optimization of query performance of columnar storage, data compression, distribution keys, as well as sort keys that minimize the I/O and also enhance response time whenever serving dashboards to multiple users at the same time.




Figure 3: Scalability and Automation Workflows

BI automation means the ability to optimize the functions of data pipelines, reporting, and alert systems without having to access them on a regular basis. When using Power BI and Redshift, it may be automated through regular refresh of the datasets, using auto ETL/ELT pipelines, REST API calls, and workflow management using Power Automate. Such automated procedures ensure that the dashboards are more or less real-time and that there are few delays between data generation and

extracting actionable insights. Furthermore, automation minimizes human error when performing routine tasks, enhances reliability, and enables businesses to be consistent in the reporting outputs. These benefits come with enormous issues that organizations should consider despite the merits of scalability and automation. In large datasets involving petabytes of data or in high-volume reporting needs, large applications in the cloud environment may be too expensive in terms of size

of resource usage, query implementation rate, and storage, and may be too costly. Moreover, it can also be doubled by sub-optimally optimized queries or not adequately optimized data models and may also incur bottlenecks, failed updates of data, and lagging insights, making BI systems undermined due to the poorly managed workflow of automation processing (e.g., scheduled refresh, API triggered, or ETL pipelines). Other problems that come about in organizations are operational problems, such as keeping data under control, compliance, and user control amongst the various

reporting levels in organizations, as well as technical problems that occur in organizations. To counter those risks, businesses have to be proactive with active monitoring and logging, a policy of governance, and cost control plans, including query scheduling, pre-aggregated tabs, incremental loads of data, and prioritization of workloads. Besides that, to improve performance and reliability, resource consumption and workflow alerting based on predictive analytics can be used.

Table 3: Scalability and Automation Considerations

Aspect	Strategy/Feature	Benefits	Challenges/Limitations
Data Scalability	Elastic cluster resizing in Redshift	Handles variable workloads efficiently	Increased cost with frequent scaling
Query Performance	Columnar storage, distribution/sort keys, compression	Faster analytics for large datasets	Requires careful schema design and optimization
Concurrency Scaling	Redshift concurrency scaling	Supports multiple simultaneous dashboard users	Can incur additional compute costs
Automated Data Pipelines	ETL/ELT pipelines using AWS Glue/Lambda	Reduces manual intervention, ensures consistent data	Requires proper monitoring and error handling
Scheduled Refresh	Power BI dataset refresh schedules	Keeps dashboards up-to-date with minimal manual effort	Overlapping refreshes can cause bottlenecks
Real-Time Data Access	DirectQuery in Power BI	Provides near real-time analytics	Can increase query latency and Redshift load
Workflow Automation	Power Automate integration	Automates alerts, notifications, and workflows	Complexity increases with multiple automated tasks
Cost Optimization	Incremental loads, query scheduling	Reduces unnecessary resource usage	Needs ongoing monitoring and tuning

IMPLEMENTATION STRATEGIES FOR BI ARCHITECTURES

Successful BI implementation relies on careful planning of ETL processes, migration strategies, and operational best practices. Key considerations include designing efficient data pipelines, ensuring data quality, minimizing downtime during migration, and maintaining compliance and

security. Challenges such as data inconsistencies, large dataset transfers, and operational bottlenecks must be addressed through robust solutions like incremental migration, validation, monitoring, and staff training. Table 3 summarizes common implementation strategies, potential issues, and recommended solutions.

Table 4: Implementation Strategies, Challenges, and Solutions in BI Deployment

Strategy Area	Description	Common Challenges	Recommended Solutions	Best Practices
ETL Design	Extract, transform, and load pipelines to integrate data from multiple sources	Data inconsistencies, complex transformations	Use automated ETL tools, and validate data before loading	Modular ETL design, version control
Migration Planning	Moving data from legacy to new BI	Downtime, large data volumes, and format	Phased migration, incremental data	Pre-migration profiling, backup

	systems	incompatibility	transfer	systems
Data Quality Management	Ensuring accuracy, completeness, and consistency	Missing, duplicate, or corrupted data	Data cleansing, validation scripts	Continuous monitoring, automated alerts
Performance Optimization	Ensure fast queries and reporting	Slow processing, bottlenecks	Indexing, parallel processing, resource scaling	Monitor performance metrics, optimize queries
Security & Compliance	Protect sensitive data during migration and operation	Unauthorized access, regulatory violations	Encryption, access control, and audit trails	Regular compliance reviews, role-based access
Automation	Reducing manual effort in data pipelines and reporting	Human errors, slow processes	Automated ETL and reporting tools	Schedule automation, monitor logs
Staff Training	Ensuring the team can operate the new BI system	Misuse of tools, errors in reporting	Training sessions, documentation	Continuous learning, knowledge sharing
Monitoring & Maintenance	Keeping the system reliable and efficient	Undetected failures, slow updates	Monitoring dashboards, alert systems	Routine maintenance, periodic audits

RESEARCH DIRECTIONS

BI systems are transforming their conventional reporting to become proactive decision-making systems. Integration of the Machine Learning (ML) models can enhance the accuracy and relevance of insights, and Non-technical users can open the data through the functions of Natural Language Processing (NLP). As organizations continue increasing the use of Artificial Intelligence (AI) and ML, the literature is increasingly committing efforts to develop model interpretability, reduce computational costs, and improve data privacy and security under clouds in the BI systems.

Among the key tendencies that led to the development of scalable BI, there must be real-time streaming analytics and multi-cloud/hybrid

models. The contemporary businesses using IoT sensors, web applications, and transaction engines are producing huge volumes of data within a short time. These dynamic streams of data are stored, processed, and analyzed in AWS Kinesis, Kafka, and Redshift streaming and incorporated into visualization platforms such as Power BI. The capability is highly valued by the financial industry, health care, and supply chain management sectors, where timely decision-making is highly valued. The hybrid and multi-cloud is scalable, and sensitive data may be stored on-premises and utilize cloud functions to undertake analytics. The comparison of real-life illustrations of scalable BI architecture implementation on the Snowflake + Tableau and BigQuery + Looker is given in Table 5 below:

Table 5: Comparative Case Study of Scalable BI Architectures

BI Architecture	Example Organization	Key Outcomes	Benefits	Challenges	Cost / Observed Benefits	Cost Comparison with Other Architectures
Snowflake + Tableau	Pharmaceutical Company (e.g., Thorogood)	Unified data warehouse, enhanced dashboard performance	Scalability, flexibility, and improved data accessibility	Integration complexity, training requirements	Reduced infrastructure costs, improved decision-making speed	Moderate Lower infrastructure cost than on-prem; slightly higher than BigQuery due to Tableau licensing

Snowflake + Tableau	Advertising Agency (e.g., BVK)	Unification of marketing data, improved client insights	Scalable analytics platform, adaptable to diverse clients	Data integration challenges	Lower data processing costs, increased profitability	Moderate to High Tableau licensing adds cost; cheaper than legacy BI stacks
BigQuery + Looker	Renewable Energy Company (e.g., GreenWind Group)	Real-time analytics, streamlined procurement insights	Cost-effective, scalable, real-time data processing	Data migration challenges, user adoption issues	Lower operational costs, faster insights delivery	Low Pay-as-you-go model, minimal infrastructure overhead
BigQuery + Looker	Retail Chain (e.g., Target)	Centralized sales and inventory analytics	Real-time reporting, enhanced supply chain efficiency	Handling high-volume transaction data	Reduced stock-outs, improved operational efficiency	Low to Moderate Economical for large-scale real-time queries
Snowflake + Power BI	Logistics Company (e.g., FedEx)	Shipment tracking, delivery performance dashboards	Improved delivery predictions, scalable reporting	Data latency, integration with legacy systems	Reduced operational costs, better customer satisfaction	Moderate to Cheaper visualization (Power BI), but Snowflake storage adds recurring cost
Power BI + Redshift	Manufacturing Enterprise (e.g., Caterpillar Inc.)	KPI dashboards for production and supply chain tracking	Tight AWS ecosystem integration, strong query performance	Limited multi-cloud flexibility, higher setup effort	Competitive total cost, good ROI for AWS-native users	Moderate to High, Lower than Snowflake + Tableau, but less flexible for multi-cloud setups

CONCLUSION

AWS Redshift and Microsoft Power BI will be suitable to offer a strong framework of automated and scalable Business Intelligence reporting. Redshift is a quick and high-scaling data warehouse, and it can process large volumes of data with its parallel processing system and columnar data storage systems, yet Power BI might be implemented to generate interactive dashboards, sophisticated analytics, and automated reporting pipelines. Combined, the two platforms can assist the organizations in developing pipelines of information that work, reduce the use of manual processes, and provide information that is timely and actionable in making good decisions. Analogs of BI architecture are drawn in a manner that on-premise BI architectures are unfavored as opposed to cloud-based BI architectures and hybrid architectures virtues of assisting scales, flexibility, and automation. Despite these benefits, it has issues to address organizational issues such as optimization of queries, cost control, and monitoring of workflow to ensure that organizations remain efficient and stable.

REFERENCES

1. Kimball, R., & Ross, M. "The data warehouse toolkit: The definitive guide to dimensional modeling." *John Wiley & Sons*, (2013).
2. Llave, M. R. "Business intelligence and analytics in small and medium-sized enterprises: A systematic literature review." *Procedia Computer Science* 121 (2017): 194-205.
3. Gupta, A., Agarwal, D., Tan, D., Kulesza, J., Pathak, R., Stefani, S., & Srinivasan, V. "Amazon redshift and the case for simpler data warehouses." *Proceedings of the 2015 ACM SIGMOD international conference on management of data*. (2015).
4. Biswas, D. "An Overview Of Data Warehousing And Olap Technology." (2020).
5. Martins, A., Martins, P., Caldeira, F., & Sá, F. "An evaluation of how big-data and data warehouses improve business intelligence decision making." *World Conference on Information Systems and Technologies*. Cham: Springer International Publishing, (2020).

6. Chaalal, H., Hamdani, M., & Belbachir, H. "Finding the best between the column store and row store Databases." *Proceedings of the 10th International Conference on Information Systems and Technologies*. (2020). Fox, G., Qiu, J., Jha, S., Ekanayake, S., & Kamburugamuve, S. "Big data, simulations and HPC convergence." *Workshop on Big Data Benchmarks*. Cham: Springer International Publishing, (2015).
7. Gaardboe, R., & Jonassen, T. S. "Business intelligence success factors: a literature review." *Journal of Information Technology Management* 29.1 (2018): 1-15.
8. Cappelli, P., & Tavis, A. "The performance management revolution." *Harvard business review* 94.10 (2016): 58-67.
9. Ogiela, L., Ogiela, M. R., & Ko, H. "Intelligent data management and security in cloud computing." *Sensors* 20.12 (2020): 3458.
10. Pherson, R. H., & Heuer Jr, R. J. "Structured analytic techniques for intelligence analysis." *Cq Press*, (2019).
11. Harris, J. G., & Craig, E. "Developing analytical leadership." *Strategic HR Review* 11.1 (2011): 25-30.
12. Armenatzoglou, N., Basu, S., Bhanoori, N., Cai, M., Chainani, N., Chinta, K., ... & Terry, D. "Amazon Redshift re-invented." *Proceedings of the 2022 International Conference on Management of Data*. (2022).
13. Zardbani, F., Afshani, P., & Karras, P. "Revisiting the Theory and Practice of Database Cracking." *EDBT*. (2020).
14. Adelusi, B. S., Ojika, F. U., & Uzoka, A. C. "A Conceptual Model for Cost-Efficient Data Warehouse Management in AWS, GCP, and Azure Environments." (2022).
15. Agrawal, D., Das, S., & El Abbadi, A. "Data Management in the Cloud." *Springer Nature*, (2022).
16. Pandis, I. "The evolution of Amazon redshift." *Proceedings of the VLDB Endowment* 14.12 (2021): 3162-3174.
17. Pavlo, A., Angulo, G., Arulraj, J., Lin, H., Lin, J., Ma, L., ... & Zhang, T. "Self-Driving Database Management Systems." *CIDR*. Vol. 4. (2017).
18. Shadbolt, N., O'Hara, K., De Roure, D., & Hall, W. "The theory and practice of social machines." *New York: Springer International Publishing*, (2019).
19. Francis, R., Gupta, R., & Oke, M. "Amazon Redshift: The Definitive Guide: Jump-Start Analytics Using Cloud Data Warehousing." *O'Reilly Media, Inc.*, (2023).
20. Bojčić, I., Marjanović, Z., Turajlić, N., Petrović, M., Vučković, M., & Jovanović, V. "A comparative analysis of data warehouse data models." *2016 6th International Conference on Computers Communications and Control (ICCCC)*. IEEE, (2016).
21. Saxena, M., Sowell, B., Alamgir, D., Bahadur, N., Bisht, B., Chandrachud, S., ... & Shah, M. A. "The story of AWS Glue." *Proceedings of the VLDB Endowment* 16.12 (2023): 3557-3569.
22. Kulkarni, A. "Amazon Redshift: Performance tuning and optimization." *International Journal of Computer Trends and Technology* 71.2 (2023): 40-44.
23. Vajpayee, A., Mohan, R., & Chilukoori, V. V. R. "Building scalable data architectures for machine learning." *International Journal of Computer Engineering and Technology (IJCET)* 15.4 (2024): 308-320.
24. Ahmadi, S. "Optimizing data warehousing performance through machine learning algorithms in the cloud." *International Journal of Science and Research* 12.12 (2023): 1859-1867.

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:

Patel, P. " Scalable and Automated Business Intelligence Reporting Using Power Bi and Aws Redshift." *Sarcouncil Journal of Engineering and Computer Sciences* 4.10 (2025): pp 386-395.