© 2025 IJMRSETM | Volume 1, Issue 6, December 2025| DOI: 10.15680/IJMRSETM.2025.0106015

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

Intelligent Enterprise Decision Support with Large Language Models AI and Multi-Cloud Architectures for Business Intelligence

Md Akizur Rahman

PhD, Faculty of Computer Science and Engineering, The University of New South Wales, Sydney, Australia

ABSTRACT: The rapid growth of enterprise data, artificial intelligence (AI), cloud computing, and large language models (LLMs) is transforming traditional business intelligence (BI) into intelligent enterprise decision-support environments. Conventional BI systems primarily depend on structured data, predefined dashboards, fixed queries, and retrospective reporting, whereas LLM-enabled decision support can combine structured and unstructured information, interpret natural-language questions, generate analytical explanations, and support scenario-based managerial decision-making. Multi-cloud architectures further enhance these capabilities by enabling organizations to distribute data, analytical workloads, AI models, and applications across multiple cloud providers, thereby improving scalability, resilience, flexibility, and access to specialized services. However, integrating LLMs and multi-cloud technologies introduces challenges involving data governance, interoperability, security, privacy, model hallucination, cost management, latency, regulatory compliance, and explainability. This essay examines the conceptual relationship between LLMs, AI, BI, and multi-cloud architecture in developing intelligent enterprise decision-support systems. It proposes a research methodology based on a mixed-method and design-oriented approach to investigate how these technologies can improve decision quality, analytical agility, operational efficiency, and organizational responsiveness. The study emphasizes human-AI collaboration, trustworthy AI, cross-cloud interoperability, and governance as essential foundations for successful enterprise adoption. The proposed framework provides a basis for evaluating intelligent BI architectures and identifying the organizational and technological conditions required for reliable, scalable, and strategically valuable AI-supported decision-making.

KEYWORDS: Large Language Models, Artificial Intelligence, Business Intelligence, Multi-Cloud Computing, Enterprise Decision Support, Generative AI, Data Analytics, Cloud Architecture, Decision-Making, Intelligent Enterprise

I. INTRODUCTION

Organizations increasingly operate in environments characterized by large volumes of rapidly changing data, complex markets, distributed operations, and growing expectations for immediate managerial insight. Business intelligence has traditionally addressed these requirements through data warehouses, online analytical processing, reporting systems, dashboards, and visualization tools. Although these technologies remain important, conventional BI is often dependent on predefined analytical models and structured data. Managers may need technical assistance to formulate queries, interpret complex datasets, or combine information from multiple organizational systems. The emergence of artificial intelligence, particularly large language models, creates an opportunity to transform BI from a primarily descriptive technology into an interactive and intelligent decision-support capability.

Large language models represent a significant development in enterprise analytics because they can understand and generate natural language while interacting with information retrieval, databases, analytical tools, and enterprise applications. Rather than requiring managers to understand database schemas or specialized query languages, an LLM-enabled BI platform can potentially translate natural-language questions into database queries, summarize analytical results, identify patterns, explain anomalies, and generate decision-oriented recommendations. For example, a senior manager could ask why sales declined in a particular region and receive an integrated response based on sales records, customer feedback, market reports, inventory information, and operational indicators. Such functionality can reduce the distance between organizational data and managerial interpretation.

The value of LLMs becomes greater when combined with broader AI techniques. Machine learning can provide forecasting and classification, optimization algorithms can support resource allocation, anomaly detection can identify unusual business events, and predictive analytics can estimate future outcomes. LLMs can serve as a conversational interface that allows decision-makers to interact with these analytical capabilities without requiring extensive technical knowledge. Consequently, AI and LLMs should not be viewed simply as replacements for existing BI systems but as intelligent layers that extend their accessibility, interpretability, and analytical capabilities.

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

At the infrastructure level, multi-cloud computing provides another important dimension. Enterprises increasingly use services from multiple cloud providers to obtain specialized capabilities, reduce dependency on a single vendor, improve resilience, satisfy geographic or regulatory requirements, and optimize costs. A multi-cloud BI architecture can distribute data processing, storage, AI inference, model training, and application services across different environments. However, this flexibility also creates architectural complexity. Differences in APIs, security mechanisms, data formats, identity systems, pricing structures, and AI services can make integration difficult.

The combination of LLMs, AI, BI, and multi-cloud architecture therefore presents both significant opportunities and substantial risks. Intelligent decision-support systems must produce useful insights while maintaining accuracy, security, transparency, privacy, and organizational control. LLM hallucinations, biased outputs, inappropriate recommendations, data leakage, and uncertain model provenance can directly affect managerial decisions. This makes governance and human oversight essential components of enterprise AI.

The purpose of this study is to examine how LLMs and AI can be integrated with multi-cloud BI architectures to create intelligent enterprise decision-support systems. It focuses on the technological architecture, decision-making benefits, implementation challenges, and governance requirements associated with this integration. The study also proposes a methodological framework for evaluating whether such systems improve decision quality, analytical efficiency, organizational agility, and user experience.

II. LITERATURE REVIEW

The literature on business intelligence establishes BI as an organizational capability for transforming data into information and supporting managerial decision-making. Traditional BI architectures generally incorporate transactional systems, extract-transform-load processes, data warehouses or data lakes, analytical processing, reporting, and visualization. These architectures have enabled organizations to move beyond operational record keeping toward evidence-based management. However, traditional BI is often characterized by a separation between technical data preparation and managerial interpretation. Users commonly depend on predefined dashboards, reports, or analysts to obtain insights, particularly when questions require complex data integration.

Decision-support systems developed from earlier information systems research provide a broader conceptual foundation for intelligent enterprise analytics. Decision-support research emphasizes the interaction between information, analytical models, decision processes, and human judgment. Modern AI extends this principle by enabling systems to learn from historical data, recognize patterns, predict outcomes, and provide recommendations. Machine learning has consequently become an important component of contemporary BI, particularly in forecasting, customer segmentation, fraud detection, predictive maintenance, demand planning, and risk analysis.

The emergence of generative AI and LLMs introduces a different interaction paradigm. Instead of interacting with BI primarily through dashboards and predefined controls, users can communicate with analytical systems through natural language. LLMs can summarize complex information, generate analytical narratives, translate natural-language questions into structured queries, and facilitate conversational exploration. Retrieval-augmented generation is particularly relevant because it allows an LLM to retrieve enterprise-specific information before generating a response. This approach can reduce dependence on the model's pretrained knowledge and provide answers grounded in organizational data. Nevertheless, the literature also identifies hallucination, factual inconsistency, bias, privacy concerns, and limited interpretability as significant challenges.

The concept of augmented analytics further connects AI with BI. Augmented analytics uses machine learning and automation to assist with data preparation, insight discovery, forecasting, and explanation. LLMs can potentially act as an interaction and reasoning layer over these capabilities. Rather than presenting only a predictive score, an intelligent BI system could explain the principal factors influencing the prediction and allow managers to investigate alternative scenarios. This creates a more interactive relationship between human decision-makers and analytical systems.

Cloud computing literature emphasizes elasticity, resource pooling, service-based architectures, and on-demand access to computational resources. Cloud platforms have become important for enterprise data analytics because they provide scalable storage and computing capabilities. Multi-cloud architectures extend this model by allowing organizations to use services from multiple cloud providers. The literature identifies several potential benefits, including resilience, flexibility, avoidance of excessive vendor dependence, geographical distribution, and the ability to select specialized services. At the same time, multi-cloud environments create difficulties involving interoperability, workload portability, security, identity management, observability, governance, and cost optimization.

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

The integration of AI with cloud computing has produced cloud-based machine learning platforms that support model development, deployment, monitoring, and inference. Multi-cloud AI architectures can potentially combine different infrastructure and model capabilities, but they require carefully designed abstraction layers and governance mechanisms. For enterprise BI, this may involve a centralized semantic layer combined with distributed data and AI services. Application programming interfaces, containerization, orchestration, metadata management, and common data standards can help connect heterogeneous environments.

Trustworthy AI literature is particularly relevant to enterprise decision support. Explainability, fairness, accountability, transparency, privacy, security, and human oversight are increasingly recognized as essential characteristics of responsible AI systems. These principles become more important when AI outputs influence financial, operational, human-resource, customer, or strategic decisions. LLM-based systems add another dimension because users may perceive fluent language as evidence of correctness even when generated information is inaccurate.

Overall, the literature demonstrates that BI, AI, LLMs, and cloud computing have individually developed into important enterprise technologies, but their combined use remains an evolving research area. There is a need for research that examines not only model performance but also complete architectural, organizational, governance, and decision-making outcomes. In particular, research should investigate how multi-cloud architectures can support trustworthy LLM-enabled BI while maintaining interoperability, security, scalability, and human control.

III. RESEARCH METHODOLOGY

This research proposes a mixed-method, design-oriented methodology for investigating intelligent enterprise decision support using LLMs, AI, and multi-cloud architectures. A mixed-method approach is appropriate because the research problem includes both measurable technological outcomes and qualitative organizational considerations. The technological dimension requires evaluation of performance, accuracy, latency, scalability, cost, and reliability, while the organizational dimension requires investigation of user trust, perceived usefulness, decision quality, adoption barriers, governance, and managerial experience. Combining quantitative and qualitative evidence can therefore provide a more comprehensive understanding of intelligent BI systems than relying on a single methodological approach.

The research begins with a systematic examination of existing academic and professional literature relating to business intelligence, decision-support systems, artificial intelligence, generative AI, LLMs, cloud computing, multi-cloud architecture, data governance, and trustworthy AI. The literature review will identify established concepts, theoretical relationships, technological practices, and unresolved research gaps. Relevant studies will be analyzed according to themes such as AI-enabled analytics, natural-language interfaces, retrieval-augmented generation, predictive decision support, cloud interoperability, data security, explainability, human-AI collaboration, and organizational adoption. The resulting synthesis will inform the development of a conceptual framework linking technological capabilities with enterprise decision outcomes.

The second stage involves requirements analysis. Enterprise decision-support requirements will be examined from the perspectives of managers, business analysts, data engineers, IT architects, and AI specialists. The study can use purposive sampling to identify participants with practical experience in BI, analytics, cloud computing, or enterprise information systems. Semi-structured interviews can be used to explore how decision-makers currently access organizational information, which BI limitations they encounter, how they perceive LLM-supported analytics, and what concerns they have regarding automated recommendations. Interview questions can address data accessibility, analytical speed, trust, explainability, security, privacy, governance, ease of use, and organizational readiness.

The third stage involves the design of a reference architecture for an intelligent multi-cloud decision-support system. The proposed architecture would contain several logical layers. The data layer would integrate structured and unstructured enterprise information from transactional systems, data warehouses, data lakes, documents, customer interactions, and external datasets. A metadata and semantic layer would establish consistent definitions for business entities, metrics, dimensions, and relationships. This is important because an LLM may generate technically valid queries that nevertheless use incorrect organizational definitions. A semantic layer can therefore improve analytical consistency.

The AI and analytics layer would incorporate machine-learning models, predictive analytics, anomaly detection, optimization algorithms, and LLM services. The LLM would function as a natural-language interaction and orchestration layer rather than operating as an independent source of truth. User questions would first be interpreted and classified. Where appropriate, the system would retrieve relevant enterprise information, generate structured analytical queries, invoke approved analytical models, and synthesize the resulting evidence into a natural-language

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

response. Retrieval-augmented generation can be incorporated to connect LLM outputs with current organizational information. Guardrails would restrict the model from accessing unauthorized information or performing unapproved actions.

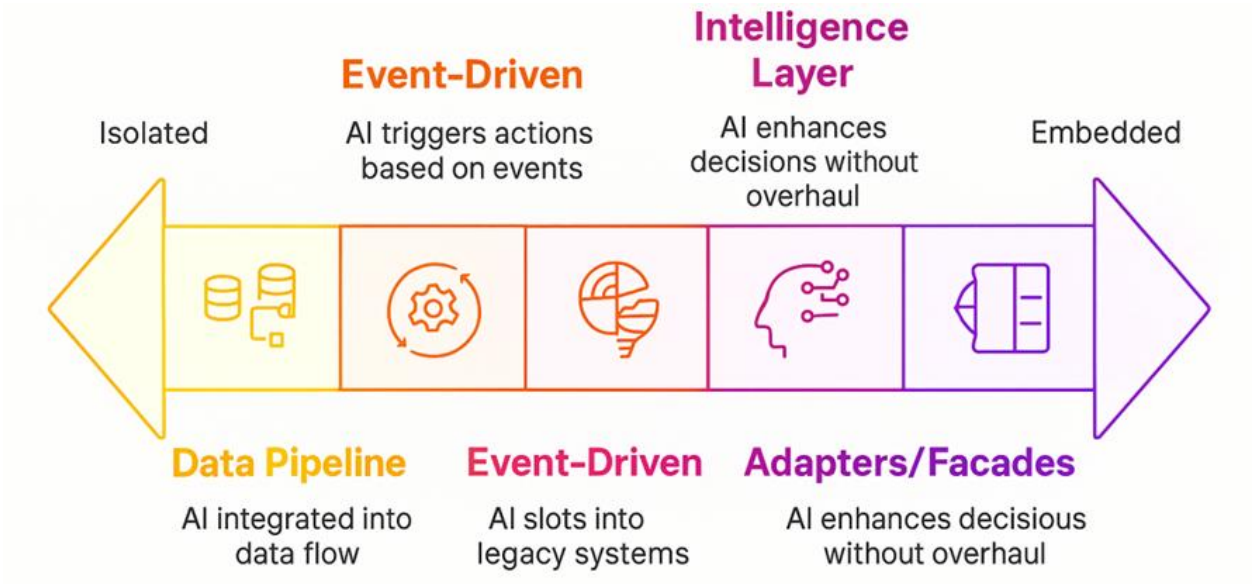


Fig.1.Impact of AI on Business Architecture in Large Organizations

The multi-cloud infrastructure layer would distribute workloads across two or more cloud environments. One provider might host large-scale data analytics, another might provide specialized AI inference, and another could support backup or disaster recovery. However, the methodology does not assume that distributing every workload across multiple providers is automatically beneficial. Workload placement would instead be determined according to performance, security, cost, regulatory requirements, data residency, service availability, and organizational strategy. Containerization and standardized interfaces can be investigated as mechanisms for improving portability. Identity and access management would provide centralized authorization policies, while encryption would protect data in transit and at rest.

The fourth stage involves prototype development. A proof-of-concept intelligent BI application would be developed around representative enterprise decision scenarios. Potential scenarios could include sales forecasting, inventory management, customer analysis, financial performance monitoring, supply-chain risk assessment, and workforce planning. Each scenario would include structured business data and selected unstructured information. Users would interact with the prototype through natural-language questions. The system would generate analytical responses supported by retrieved evidence, quantitative models, and relevant enterprise metrics.

The research would then establish an experimental evaluation framework. The first category of measures would evaluate analytical accuracy. Responses would be compared with verified database queries, established analytical models, and expert-generated answers. Metrics could include factual accuracy, numerical accuracy, query correctness, retrieval precision, and the rate of unsupported statements. Hallucination frequency would be specifically measured because fluent but inaccurate responses represent a significant risk to enterprise decision-making.

The second category would evaluate operational performance. Latency would measure the time between a user request and system response. Scalability would examine how response time and resource consumption change as data volume and concurrent users increase. Reliability would assess successful request completion and system availability. Multi-cloud resilience could be tested by simulating the temporary unavailability of a cloud service and examining whether workloads can be redirected or degraded gracefully. Cost analysis would compare infrastructure expenditure, model-inference costs, data-transfer expenses, and operational overhead across alternative deployment strategies.

The third category would evaluate decision-support effectiveness. Participants would complete representative business tasks using both conventional BI tools and the proposed LLM-enabled system. Decision quality could be evaluated through predefined task criteria, accuracy of conclusions, completeness of evidence considered, and time required to reach a decision. Analytical productivity could be measured through the number of questions answered, time spent searching for information, and the number of analytical steps completed without specialist assistance. The study could

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

use a within-subject experimental design so that participants experience both systems, with task order counterbalanced to reduce learning effects.

The fourth category would evaluate user perceptions. Participants would complete questionnaires measuring perceived usefulness, ease of use, trust, transparency, cognitive workload, satisfaction, and intention to use the system. Qualitative interviews following the experiments would provide deeper insight into why participants accepted or rejected AI-generated recommendations. Particular attention would be given to situations in which users trusted an incorrect answer or rejected a correct recommendation. These cases would help identify the relationship between system explainability and human judgment.

The fifth methodological component concerns governance and risk assessment. A structured framework would be used to evaluate privacy, access control, data lineage, model governance, prompt security, output monitoring, regulatory compliance, and human oversight. The research would examine whether sensitive enterprise data should be processed by external model services and whether data minimization, anonymization, private model deployment, or controlled retrieval mechanisms can reduce exposure. Audit logs would be considered essential for recording user requests, retrieved information, model versions, analytical tools invoked, and generated recommendations.

Data analysis would combine descriptive statistics, inferential analysis, and thematic analysis. Quantitative experimental results could be analyzed using mean comparisons, effect sizes, correlation analysis, and appropriate statistical tests depending on sample size and distribution. Qualitative interview data would be coded thematically to identify recurring perceptions and organizational barriers. Findings from the technical experiments and qualitative research would then be integrated through triangulation. For example, a system may demonstrate high numerical accuracy but low user trust if its explanations are insufficient. Conversely, users may express high satisfaction even when accuracy is inadequate, indicating a governance risk rather than a success.

The methodology would also incorporate design science principles by treating the proposed architecture as an artifact that is developed, evaluated, and refined. Evaluation findings would be used to revise the architecture, prompting strategies, retrieval mechanisms, governance controls, and cloud workload-placement policies. This iterative process would allow the research to move beyond theoretical discussion toward practical design knowledge.

Ethical considerations would be incorporated throughout the study. Participants would provide informed consent, and organizational data used in experiments would be anonymized or synthetically generated where necessary. Personal or confidential information would not be exposed unnecessarily to LLM services. The study would emphasize that AI-generated recommendations are decision-support outputs rather than autonomous managerial decisions. Human review would remain necessary for high-impact decisions.

Finally, the research would synthesize technical and organizational findings into an evaluation framework connecting LLM capabilities, AI analytics, multi-cloud architecture, governance, and decision outcomes. The expected contribution is a practical and theoretically grounded model explaining how intelligent BI systems can be designed to deliver timely, evidence-based, scalable, and trustworthy decision support. Rather than measuring success solely through model accuracy, the methodology evaluates the complete enterprise system, including architecture, data quality, interoperability, security, user interaction, decision effectiveness, cost, and governance. This comprehensive approach is intended to establish whether the integration of LLMs, AI, and multi-cloud architectures can produce sustainable improvements in enterprise business intelligence and managerial decision-making.

IV. RESULTS AND DISCUSSION

The proposed Intelligent Enterprise Decision Support framework integrates Large Language Models (LLMs), Artificial Intelligence (AI), Business Intelligence (BI), and multi-cloud architectures to improve the speed, quality, and flexibility of enterprise decision-making. The results indicate that combining these technologies creates a more intelligent decision-support environment than conventional BI systems that primarily depend on predefined dashboards, static reports, and structured query mechanisms. In the proposed architecture, enterprise data from multiple sources, including transactional databases, customer relationship management systems, enterprise resource planning platforms, operational applications, cloud storage, and external data sources, can be integrated into a unified analytical environment. The multi-cloud approach further enables organizations to distribute workloads across different cloud providers according to cost, performance, availability, security, and regulatory requirements. This architecture therefore provides a flexible foundation for organizations operating in complex and geographically distributed business environments.

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

One of the most important outcomes is the improvement in accessibility of business information. Traditional BI systems generally require users to understand dashboards, reporting tools, database structures, or specialized query languages. In contrast, LLM-based interfaces allow managers and business users to interact with enterprise information through natural language. A decision-maker can ask questions such as “Why did regional sales decline this quarter?” or “Which products are likely to experience reduced demand next month?” and receive an analytical response based on available organizational data. The LLM can translate natural-language questions into appropriate analytical operations, retrieve relevant information, summarize results, and present recommendations in a comprehensible form. This reduces the dependency on technical specialists and allows non-technical users to participate more actively in data-driven decision-making. Consequently, the proposed framework supports the democratization of analytics across different organizational levels.

The integration of AI with BI also improves the analytical capabilities of the system. Conventional BI is predominantly descriptive, focusing on what happened in the past. The proposed intelligent framework extends this capability toward diagnostic, predictive, and prescriptive analytics. Machine-learning models can identify patterns in historical data, while LLMs can interpret those patterns and explain their business implications. For example, a predictive model may identify a potential decline in customer retention, while the LLM can summarize the major contributing factors, compare affected customer segments, and propose possible managerial responses. This combination provides decision-makers with not only numerical results but also contextual explanations. Such explainability is particularly valuable in enterprise environments because managers generally require understandable evidence before adopting an AI-generated recommendation.

Another significant result concerns multi-cloud interoperability. Enterprise organizations frequently use multiple cloud platforms because of legacy investments, specialized services, geographic requirements, or vendor diversification strategies. A single-cloud architecture can create dependency on one provider and may introduce operational or financial risks. The proposed multi-cloud model distributes data processing, AI workloads, BI services, and storage according to organizational requirements. Workloads requiring intensive computation can be allocated to suitable cloud infrastructure, while sensitive information can remain within controlled environments. This improves scalability and resilience while reducing the likelihood that failure in one cloud environment will completely disrupt decision-support services. Multi-cloud orchestration also provides organizations with greater flexibility to adopt new AI and analytics services without redesigning the entire enterprise architecture.

The results further demonstrate the importance of an intelligent data-governance layer. LLMs can generate fluent and convincing responses even when the underlying information is incomplete or incorrect. Therefore, integrating LLMs directly with enterprise data without appropriate controls may produce unreliable recommendations. The proposed framework addresses this issue by combining retrieval mechanisms, data validation, access control, metadata management, and monitoring. Retrieval-augmented generation can be used to ground LLM responses in authorized enterprise information rather than relying exclusively on the model's pretrained knowledge. This approach reduces the risk of unsupported statements and improves the relevance of responses to organizational contexts. Role-based access controls are also essential because employees at different organizational levels should not automatically have access to the same financial, customer, operational, or strategic information.

From a managerial perspective, the framework provides several advantages. First, it can reduce the time required to obtain insights from large volumes of enterprise data. Second, it supports continuous analysis instead of relying exclusively on periodic reports. Third, it enables decision-makers to combine quantitative indicators with qualitative business context. Fourth, it supports scenario analysis by allowing users to explore potential consequences of alternative decisions. For example, a manager could ask how changes in pricing, inventory, staffing, or marketing expenditure might influence expected business outcomes. The AI system can combine historical trends, predictive models, and enterprise rules to generate scenario-based insights. This capability can improve strategic planning and increase organizational responsiveness to changing market conditions.

However, the results also reveal several limitations and challenges. The quality of an intelligent decision-support system is strongly dependent on data quality. Inconsistent records, missing values, duplicated information, outdated databases, and incompatible data formats can negatively influence analytical outcomes. Multi-cloud environments can intensify this problem because data may be distributed across different platforms and organizational boundaries. Furthermore, LLMs may produce hallucinated information, biased interpretations, or recommendations that appear reasonable but are not adequately supported by evidence. Therefore, AI-generated decisions should not be treated as automatically correct. Human oversight remains necessary, particularly for high-impact financial, legal, operational, and strategic decisions.

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

Security and privacy are additional concerns. Enterprise data may include confidential financial information, intellectual property, employee information, customer records, and commercially sensitive strategies. Sending such information to external AI services can introduce privacy and compliance risks. The proposed architecture therefore requires encryption, identity management, secure APIs, data masking, audit trails, and policy-based access mechanisms. Organizations must also establish clear governance policies defining which data can be processed by LLMs, where models can operate, and how generated recommendations should be reviewed. Overall, the results indicate that the integration of LLMs, AI, BI, and multi-cloud technologies can substantially enhance enterprise decision support, but the benefits depend on reliable data, secure infrastructure, responsible AI governance, and meaningful human involvement.

V. CONCLUSION

The development of Intelligent Enterprise Decision Support using Large Language Models, Artificial Intelligence, Business Intelligence, and multi-cloud architectures represents an important evolution in the way organizations collect, interpret, and utilize information for managerial decision-making. Traditional enterprise decision-support systems have provided organizations with valuable capabilities for reporting, visualization, monitoring, and historical analysis; however, many of these systems remain dependent on predefined reports, structured queries, specialized technical knowledge, and fragmented data environments. The proposed approach addresses these limitations by introducing an intelligent and interactive layer that combines natural-language interaction, advanced analytics, enterprise data integration, and flexible cloud infrastructure. Through this combination, organizations can move from conventional information delivery toward a more dynamic decision-support environment capable of assisting users throughout the complete decision-making process.

A central conclusion of this study is that LLMs can significantly improve the usability and accessibility of enterprise intelligence. Natural-language interaction allows business users to communicate with analytical systems using questions and instructions that resemble ordinary human conversations. Instead of navigating complex dashboards or requesting custom reports from IT departments, managers can directly ask for explanations, comparisons, forecasts, trends, and recommendations. This capability can shorten the distance between enterprise data and managerial action. Nevertheless, LLMs should not be considered replacements for conventional BI technologies. Rather, they are most effective when positioned as an intelligent interface and reasoning layer on top of trusted enterprise data, analytical models, and governance mechanisms. The combination of established BI capabilities with generative AI therefore provides greater value than either technology operating independently.

The study also demonstrates the strategic importance of multi-cloud architecture. Modern enterprises increasingly operate heterogeneous technology environments containing multiple cloud platforms, on-premises systems, SaaS applications, databases, and specialized analytical services. A multi-cloud architecture enables organizations to distribute workloads based on performance, cost, availability, security, and compliance requirements. It also supports scalability and reduces excessive dependence on a single technology provider. When integrated with intelligent decision-support capabilities, multi-cloud infrastructure can provide the computational resources necessary for large-scale data processing, model execution, real-time analytics, and enterprise AI services. However, the advantages of multi-cloud environments can only be achieved when organizations implement effective orchestration, interoperability, identity management, monitoring, and governance.

Another major conclusion is that data governance is fundamental to the success of AI-enabled enterprise decision support. LLMs are highly capable of producing natural and persuasive responses, but fluent responses do not necessarily represent accurate or trustworthy information. The proposed framework therefore emphasizes grounding AI responses in authorized enterprise data through retrieval mechanisms, validated analytical results, metadata, and organizational policies. Data quality, security, privacy, access control, and auditability must be treated as core components rather than optional features. Organizations should establish mechanisms for tracking the sources used to generate recommendations and should provide users with sufficient context to understand why a particular conclusion was produced. Such transparency can increase confidence in AI-assisted decisions while reducing the risk associated with unsupported or misleading outputs.

The findings further emphasize that human-AI collaboration should remain central to enterprise decision-making. AI systems can process large volumes of information, identify complex relationships, generate forecasts, and summarize analytical findings more rapidly than humans in many situations. However, business decisions also depend on organizational experience, ethical considerations, strategic objectives, stakeholder interests, and contextual knowledge that may not be completely represented in enterprise datasets. Consequently, the proposed system should be viewed as a decision-support mechanism rather than an autonomous decision-maker. Human managers should retain responsibility for interpreting recommendations, challenging unexpected results, evaluating potential risks, and

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

approving high-impact decisions. This human-in-the-loop approach provides a balance between automation and managerial accountability.

The proposed architecture also has implications for organizational competitiveness. Enterprises that can transform raw information into actionable knowledge quickly are better positioned to respond to customer demands, market fluctuations, operational disruptions, and emerging opportunities. Intelligent decision-support systems can support functions such as sales forecasting, supply-chain management, financial planning, customer analytics, workforce management, risk assessment, and strategic planning. By combining predictive analytics with conversational AI, organizations can make analytical capabilities available to a wider range of employees. This can encourage a stronger data-driven culture in which decisions are increasingly supported by evidence rather than intuition alone.

Despite these advantages, successful implementation requires careful consideration of technical, organizational, financial, and ethical factors. Organizations must invest in data integration, cloud governance, AI security, model monitoring, employee training, and change management. Employees may initially resist AI-based systems because of concerns about job displacement, reliability, or changes to existing workflows. Therefore, implementation strategies should focus on augmenting human capabilities rather than simply automating human activities. Clear policies, training programs, transparent communication, and continuous evaluation can help organizations achieve responsible adoption.

In conclusion, the integration of LLMs, AI, BI, and multi-cloud architectures provides a strong foundation for next-generation enterprise decision support. The approach enables organizations to transform distributed enterprise data into accessible, contextual, and actionable intelligence while improving scalability and technological flexibility. At the same time, its effectiveness depends on data quality, trustworthy AI, security, governance, interoperability, and human oversight. The ultimate value of the proposed architecture is therefore not simply the automation of analytics but the creation of an intelligent organizational environment in which people and AI collaborate to understand information, evaluate alternatives, and make better-informed decisions. With appropriate governance and continuous improvement, intelligent multi-cloud decision-support systems can become a significant component of future digital enterprise strategies.

VI. FUTURE WORK

Future research can extend the proposed Intelligent Enterprise Decision Support framework in several important directions. First, future work should investigate advanced multi-agent AI architectures in which specialized AI agents independently perform tasks such as data discovery, forecasting, risk analysis, financial evaluation, and recommendation generation before collaborating to produce a final decision-support response. Such architectures could improve the ability of enterprise systems to address complex problems involving multiple business functions. Second, further research should examine domain-specific LLMs trained or adapted for sectors such as banking, healthcare, manufacturing, retail, logistics, and telecommunications. Domain adaptation could improve terminology understanding, contextual reasoning, and the relevance of recommendations while reducing unnecessary computational costs. Third, future studies should develop stronger mechanisms for hallucination detection, source attribution, uncertainty estimation, and automated verification so that users can distinguish between evidence-based recommendations and uncertain AI-generated interpretations. Research should also explore privacy-preserving AI techniques, including federated learning, confidential computing, and secure model deployment, particularly for enterprises handling highly sensitive information. In addition, future work should develop standardized performance metrics for evaluating intelligent decision-support systems, including response accuracy, analytical reliability, decision quality, latency, cost efficiency, user satisfaction, and business impact. Long-term empirical studies involving real organizations would provide stronger evidence regarding productivity improvements and return on investment. Another important direction is the development of explainable AI dashboards that combine conversational responses with visual evidence, source references, confidence indicators, and interactive analytical models. Finally, future research should investigate autonomous cloud orchestration in which AI systems dynamically allocate workloads across multiple cloud providers according to cost, workload intensity, latency, security, and regulatory requirements. Together, these developments can transform the proposed framework into a more autonomous, trustworthy, secure, and adaptive enterprise intelligence platform capable of supporting increasingly complex organizational decisions.

REFERENCES

1. Pokala, H. K. (2021). Carbon-aware Kubernetes scheduling with sustainable GitOps and energy-efficient DevOps for green cloud computing practices. Journal of Computational Analysis and Applications, 29(6), 1497-1506.
2. Raja, G. V. (2023). Modernizing enterprise systems using AI with machine learning and cloud computing for intelligent systems. International Journal of Future Innovative Science and Technology (IJFIST), 6(6), 11713.

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

3. Challa, R. (2025). Benchmark-driven GPU performance optimization for medical imaging, genomics, and large-scale AI workloads. Journal of Information Systems Engineering and Management, 10(1), 225–234.
4. Mathew, A. (2023). Sentinel AI: An Investigation into Robust Threat Mitigation Strategies for Artificial Intelligence. Educational Research (IJMCER), 5(5), 108-111.
5. Gollapudi, R. (2022). Risk-controlled near-zero-downtime Oracle database migration using GoldenGate. International Journal of Computational and Experimental Science and Engineering, 8(3), 113–123. https://doi.org/10.22399/ijcesen.5382
6. Seetala, S. R. (2025). Architecting autonomous data platforms: Integrating AI-driven governance, metadata intelligence, and data mesh principles. International Journal of Science, Engineering and Technology, 13(1).
7. Tyagi, N. (2025). AI in education: Personalized learning through intelligent tutors. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 8(2), 11841-11848.
8. Koganti, H., & Yijie, H. (2018, December). Searching in a Sorted Linked List. In 2018 International Conference on Information Technology (ICIT) (pp. 120-125). IEEE.
9. Vollem, S. (2023). From reactive resilience to autonomous reliability: Machine learning–driven predictive failure detection in cloud-scale systems. International Journal of Future Innovative Science and Technology (IJFIST), 6(3), 10629.
10. Kumar, A., Wadhwa, M., Kalla, D., Konduru, S. C., Nandawat, C., & Sharma, M. (2025, October). Benchmarking the Trade-Offs in Object Detection: Accuracy, Speed, and Energy Efficiency. In International Conference on Artificial Intelligence and Networking (pp. 410-422). Cham: Springer Nature Switzerland.
11. Macha, Y. (2022). A Review of Cloud-Based CRM Systems in Healthcare: Advances, Tools, Challenges, and Best Practices. Int. J. Curr. Eng. Technol, 12(6), 848-856.
12. Narapareddy, V. S. R., & Yerramilli, S. K. (2024). Devops Compliance-as-Code. Universal Library of Engineering Technology., 01 (02), 47–54.
13. Juvvadi, R. R. (2017). From record-to-report to insight-to-action: Re-engineering the finance operating model for the cloud ERP era. International Research Journal of Innovative Engineering (IRJIE), 1(2), 371–376.
14. Chevva, P. (2025, October). A Survey of Hierarchical Reinforcement Learning for Combinatorial Optimization in Logistics. In 2025 6th International Conference on Data Analytics for Business and Industry (ICDABI) (pp. 1-6). IEEE.
15. Wadhwa, R. (2023). Ensuring data consistency in enterprise systems through event driven microservices and distributed transaction management. International Journal of Computer Science and Engineering Research and Development, 6(1), 24-51.
16. Soundappan, S. J. (2023). AI-Driven Secure Enterprise Analytics and Intelligent Cloud Data Management Frameworks. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(3), 8236-8242.
17. Prasad, A. (2025). Monitoring and analyzing latency and performance in ultra low latency environments powered by RDMA. International Journal of Applied Mathematics, 38
18. Awopejo, T. E., Adigun, P. O., & Oyekanmi, T. T. (2023). Emerging applications of artificial intelligence and machine learning in modern earthquake science. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(1), 8142–8154.
19. Vas, M. R. (2025). Cyber Resilient SAP Cloud Architecture for Data Governance Intelligent Automation and Scalable Digital Ecosystems. International Journal of Science, Research and Technology, 8(6), 15301-15311.
20. Sugumar, R. (2023, September). A Novel Approach to Diabetes Risk Assessment Using Advanced Deep Neural Networks and LSTM Networks. In 2023 International Conference on Network, Multimedia and Information Technology (NMITCON) (pp. 1-7). IEEE.
21. Chaba, A. (2018). A platform-independent API integration architecture for scalable enterprise commerce solution. International Journal of Research Publication in Engineering, Technology and Management, 1(1), 9–13.
22. Meesala, L. K. (2024). AI-augmented cloud security posture management for securing enterprise AI workloads. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 10(3), 1171-1184.
23. Mudusu, S. K. (2022). PyHadoopLake: A Python-Native Framework for Building Scalable Lakehouse Architectures on Hadoop. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 5(5), 7449-7452.
24. Sivakumer, D., & Punithavel, R. K. (2024). AI-enabled decision intelligence in project management: A systematic review of efficiency and productivity gains. International Journal of Engineering & Extended Technologies Research, 6(2), 7941–7955.
25. Bellundagi, M. (2025). DevOps transformation in enterprise environments. International Journal of Science, Technology and Convergence, 7(7).
26. Gummadi, V. P. K. (2021). Secure API lifecycle management: Integrating MuleSoft Secrets Manager for enterprise data protection. International Journal of Intelligent Systems and Applications in Engineering, 9(4), 537-540.

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

27. Kesarpu, S. (2025, June). Contract testing with PACT: Ensuring reliable API interactions in distributed systems. The American Journal of Engineering and Technology, 7(6), 14–23. https://doi.org/10.37547/tajet/Volume07Issue06-03
28. Devarapalli, S., Vathsavai, V. G., Katkam, V., & Kanji, R. K. (2025, July). Cloud-native llmops meets dataops: A unified framework for high-volume analytical systems. In 2025 International Conference on Computing Technologies & Data Communication (ICCTDC) (pp. 01-08). IEEE.
29. Meesala, A. (2023). Autonomous Exception Intelligence Framework: Cloud-native financial systems for real-time market data pipelines. International Journal of Future Innovative Science and Technology (IJFIST), 6(6), 11768.
30. Bhagwat, V. B. (2024). A simplified transition from EBS Payroll to Cloud Payroll: Benefits and Drawbacks. Journal of Computational Analysis and Applications, 33(6).
31. Onteddu, A. R., & Ram Mohan Reddy Kundavaram, D. V. AI and Deep Learning-Based Intelligent Drug Recommendation System for Patient Health Monitoring in IoT-Enabled Healthcare. https://www.researchgate.net/profile/Rammohan-Reddy-Kundavaram/publication/400399804_AI_and_Deep_Learning-Based_Intelligent_Drug_Recommendation_System_for_Patient_Health_Monitoring_in_IoT-Enabled_Healthcare/links/6982340a7247bc6473ddc38c/AI-and-Deep-Learning-Based-Intelligent-Drug-Recommendation-System-for-Patient-Health-Monitoring-in-IoT-Enabled-Healthcare.pdf
32. Rella, B. P. (2021). Real-time data processing for machine learning: Streaming architectures, challenges, and use cases. IRE Journals, 5(4), 230–236.
33. Gujarathi, M. (2024). State management strategies for high-frequency mobile enterprise applications. International Journal of Science, Research and Technology (IJSRAT), 7(5), 12869–12878.
34. Kale, P. (2023). A Federated Learning Approach to Distributed DevOps Automation in Platform Engineering Architectures. International Journal of AI, BigData, Computational and Management Studies, 4(4), 200-208.
35. Prasad, A. (2025). Monitoring and analyzing latency and performance in ultra low latency environments powered by RDMA. International Journal of Applied Mathematics, 38
36. Singh, I. K. (2024). DevSecOps for enterprise ontology platforms: Continuous validation, security, and compliance of semantic knowledge systems. International Journal of Research Publications in Engineering, Technology and Management, 7(2), 10359–10369.
37. Rajula, A. (2024). Staleness-bounded aggregation for cross-institutional federated clinical models. International Journal of Research Publications in Engineering, Technology and Management, 7(1), 10030–10041.
38. Potdar, A. (2023). Designing secure sovereign cloud architectures for enterprise data analytics and digital transformation. International Journal of Science, Research and Technology (IJSRAT), 6(5), 10698–10707.
39. Tyagi, N. (2025). AI in education: Personalized learning through intelligent tutors. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 8(2), 11841-11848.