International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(21s) 2026 pp. 1865-1885
https://doi.org/10.70917/ijcisim-2026-5942
Article

RETRIEVAL-AUGMENTED GENERATION ARCHITECTURE FOR AUTOMATED DECISION-MAKING AND STRATEGIC AUDITING IN CORPORATE INFORMATION SYSTEMS

Saud Ilahi[1], Dr. J. Bamini[2], Aparna Ponnuru[3], Prof. Nisha Singh[4], Vineet Kumar Mittal[5], Dr. D. Arivazhagan[6]

[1] Assistant Professor, Management and Marketing, Jazan University, Jazan, KSA.
[2] Cluster Head - UG SF Management programs, PSGR Krishnammal College for Women Peelamedu, Coimbatore.
[3] Sr. Assistant Professor, Department of CSE (AI&ML), Geethanjali College of Engineering and Technology.
[4] Professor, Department of English and Other Foreign Languages, Mahatma Gandhi Kashi Vidyapith Varanasi.
Special mention: English language correction for this Manuscript.
[5] Independent Researcher.
[6] AMET Business School, Academy of Maritime Education and Training Deemed to be University.

Abstract: The rapid adoption of Generative Artificial Intelligence (GenAI) in corporate information systems has created new opportunities for automated decision-making and strategic auditing. However, conventional Large Language Models (LLMs) may generate inaccurate, outdated, or unsupported information, creating concerns regarding reliability, transparency, accountability, and evidence verification. This study proposes a trustworthy Retrieval-Augmented Generation (RAG) architecture that integrates enterprise information retrieval, evidence validation, provenance tracking, explainability, LLM reasoning, and human oversight to support evidence-based corporate decision-making and strategic auditing. The study adopts a design science and quantitative research approach and examines the perceptions of 250 respondents with diverse professional expertise in finance, auditing, information systems, management, and risk management. The proposed framework considers RAG Architecture Characteristics, Evidence Traceability, Explainability, Human Oversight, Automated Decision-Making Effectiveness, and Strategic Auditing Effectiveness as key constructs. The findings indicate high perceptions of RAG Architecture Characteristics (Mean = 4.12), Evidence Traceability (Mean = 4.18), Explainability (Mean = 4.12), and Automated Decision-Making Effectiveness (Mean = 4.21). Human Oversight recorded the highest mean (4.42), highlighting the importance of expert validation, accountability, and intervention in high-impact AI-supported decisions. Strategic Auditing Effectiveness recorded a mean of 4.23, while Hallucination Reduction also achieved a mean of 4.23. Furthermore, RAG Architecture Characteristics demonstrated significant positive relationships with decision outcomes, with Decision Relevance showing the strongest relationship ($\beta = 0.68$, $R^2 = 0.520$, $p < 0.001$). The study concludes that evidence-grounded, human-supervised RAG can strengthen reliable, explainable, and auditable enterprise AI systems.

Keywords: Retrieval-Augmented Generation; Generative AI; Large Language Models; Corporate Information Systems; Automated Decision-Making; Strategic Auditing; AI Governance; Explainable AI; Enterprise AI

1. Introduction

Corporate information systems have become increasingly data-intensive, generating vast volumes of structured and unstructured information through Enterprise Resource Planning (ERP), Customer Relationship Management

Copyright: © 2026 by the authors

(CRM), accounting and financial systems, supply-chain platforms, human resource applications, contracts, policies, audit reports and regulatory repositories. The ability to effectively analyse these diverse information sources is critical for organisations seeking to improve strategic decision-making, risk management, regulatory compliance and internal auditing. Recent advances in Generative Artificial Intelligence (GenAI), particularly Large Language Models (LLMs), have created new opportunities to automate information analysis and provide managerial insights. However, conventional LLM-based systems primarily depend on their pre-trained knowledge and may generate plausible but inaccurate or unsupported outputs, commonly referred to as hallucinations. Such limitations create significant concerns when AI is applied to corporate decisions where accuracy, accountability, transparency and evidence verification are essential.

Retrieval-Augmented Generation (RAG) has emerged as a promising approach for addressing these limitations by connecting LLMs with external knowledge repositories. Rather than relying exclusively on pre-trained knowledge, RAG retrieves relevant and current information from enterprise databases, documents and knowledge repositories and incorporates this evidence into the generation process. This enables AI systems to produce responses that are more contextually relevant and potentially more reliable. In corporate environments, RAG can provide access to organisation-specific information while supporting source attribution and evidence-based recommendations. Consequently, RAG has significant potential for applications involving financial analysis, regulatory compliance, risk assessment, strategic planning and management decision support.

Recent developments have shifted RAG research from conventional question-answering systems towards enterprise-oriented architectures that emphasise provenance, evidence validation and controlled decision processes. Evidence provenance is particularly important in corporate applications because decision-makers need to understand the sources underlying AI-generated recommendations. Similarly, the ability to trace retrieved information and establish an audit trail can strengthen accountability and facilitate subsequent verification of AI-supported decisions.

The application of RAG in auditing represents another emerging research direction. AI-assisted auditing can potentially automate document analysis, compliance verification, anomaly identification, risk assessment and generation of preliminary audit findings. Emerging approaches have explored the integration of RAG with AI agents, workflow orchestration, query classification and multi-agent systems. Other studies have proposed specialised architectures, including Graph-RAG and dual-channel RAG, to improve retrieval and reasoning when auditing requires analysis of complex relationships among financial, organisational and regulatory information.

Despite these developments, existing research remains fragmented across RAG architecture, enterprise decision support and AI-assisted auditing. Limited research has empirically examined how different RAG architectures influence decision accuracy, evidence traceability, hallucination reduction, explainability and audit effectiveness within a unified corporate information-system environment. Therefore, an important research opportunity exists to develop and validate an integrated RAG architecture that connects enterprise information retrieval with automated decision-making and strategic auditing. Addressing this gap can contribute to the development of trustworthy, explainable and auditable enterprise AI systems capable of supporting evidence-based corporate decision-making.

2. Review of Literature

2.1 Large Language Models and Corporate Information Systems

The emergence of Large Language Models (LLMs) has created substantial opportunities for organisations to automate information processing, knowledge management and decision-support activities. LLMs can analyse large volumes of textual information and generate human-like responses, making them increasingly relevant to accounting, auditing, finance and enterprise information systems. However, their effectiveness in professional environments is constrained by hallucinations, outdated knowledge, limited domain specificity and difficulties in tracing the basis of generated responses. These limitations are particularly important in corporate decision-making, where recommendations should be supported by reliable organisational and regulatory information.

Research examining ChatGPT in accounting and auditing demonstrates both the potential and limitations of LLMs. Eulerich et al. (2024) found that improvements in model capability, additional training and access to external tools substantially enhanced performance on accounting and auditing-related tasks. Their findings suggest that LLM performance should not be considered solely in terms of the underlying model but also in terms of the information, tools and contextual resources available to the system.

This provides an important foundation for investigating RAG-based corporate information systems, where external organisational knowledge can be integrated into the LLM generation process.

2.2 Retrieval-Augmented Generation

Retrieval-Augmented Generation was initially introduced as a framework that combines a neural language generator with an external retriever to provide access to knowledge beyond the model's parametric memory (Lewis et al., 2020). The approach is particularly relevant for knowledge-intensive applications because retrieved information can provide current and domain-specific evidence to the generation process.

Subsequent research has identified RAG as an important approach for addressing several limitations of LLMs, particularly hallucination, knowledge updating and domain adaptation. Gao et al. (2024) classify RAG development into Naïve RAG, Advanced RAG and Modular RAG**,** highlighting improvements in retrieval, augmentation and generation mechanisms. They also identify evaluation, reliability and retrieval quality as important areas for further research.

Similarly, Wu et al. (2024) emphasise that RAG addresses the limitations of static LLM knowledge by retrieving information from external databases. Their survey identifies advances in retriever design, retrieval fusion, datastore updating and industrial applications, demonstrating that RAG has evolved from a basic retrieval-generation pipeline into a broader architectural paradigm.

Gupta et al. (2024) further observe that RAG has expanded across question answering, summarisation and knowledge-intensive applications while raising concerns concerning scalability, bias, robustness and ethical deployment.

For corporate information systems, these developments suggest that RAG can potentially connect LLMs with enterprise-specific information, thereby improving contextual relevance and reducing dependence on static model knowledge.

2.3 RAG Architecture and Evidence Retrieval

The effectiveness of a RAG system depends substantially on the quality of the retrieval process. Huang and Huang (2024) conceptualise RAG through four stages: pre-retrieval, retrieval, post-retrieval and generation. This perspective demonstrates that RAG performance is not determined solely by the LLM but by the interaction between information preparation, retrieval, evidence processing and generation.

For enterprise applications, this is particularly significant because organisational information is heterogeneous. Financial reports, policies, contracts, audit documents and regulatory materials differ in structure, terminology and relevance. Consequently, enterprise RAG systems need mechanisms for document classification, semantic search, metadata management, re-ranking and evidence selection.

Zhao et al. (2024) similarly argue that RAG can improve the accuracy and robustness of AI-generated content by introducing relevant information retrieval into the generation process. Their review also highlights challenges associated with data leakage, long-tail information, computational costs and system evaluation.

Thus, retrieval quality represents a fundamental determinant of the reliability of RAG-based corporate decision-support systems.

2.4 Graph-RAG and Complex Corporate Knowledge

Conventional vector-based retrieval is effective when relevant information is contained within individual documents or passages. However, strategic corporate decisions and auditing frequently require understanding relationships among multiple entities, transactions and documents**.** Graph-based approaches have therefore emerged as an important development in RAG architectures.

Peng et al. (2024) explain that GraphRAG incorporates structural relationships among entities to improve retrieval and context-aware generation. Their review identifies three major components: graph-based indexing, graph-guided retrieval and graph-enhanced generation**.**

Microsoft Research similarly describes GraphRAG as an approach combining text extraction, network analysis and LLM-based summarisation to support richer understanding of complex text collections.

This architecture is particularly relevant to auditing because audit investigations may require connections between transactions, suppliers, employees, accounts, policies and regulatory requirements. Therefore, Graph-RAG may provide advantages over conventional RAG for complex strategic auditing tasks.

2.5 Generative AI and Auditing

The accounting and auditing literature increasingly recognises the potential of GenAI to transform audit procedures. Gu et al. (2024) introduce the concept of AI co-piloted auditing, examining how foundation models and LLMs can support diverse auditing tasks. Their work indicates that GenAI can potentially assist auditors in analysing documents and performing complex audit procedures, while also highlighting the importance of appropriate model configuration and reasoning mechanisms.

More recently, Xiong et al. (2026) proposed an AI-agent framework specifically designed for auditing workflows. The framework integrates RAG with private knowledge bases, workflow orchestration, query classification and multi-agent coordination. Importantly, their proposed evaluation compares conventional RAG and Graph-RAG for auditing tasks using measures of accuracy, time and token consumption.

This research provides strong support for the relevance of RAG to auditing. However, it primarily focuses on AI-agent-based audit workflows, leaving scope for broader investigation of RAG as an integrated architecture for corporate decision-making and strategic auditing.

2.6 AI Governance, Explainability and Auditability

The increasing use of LLMs in organisational decision-making raises important questions concerning transparency, accountability and governance. Mökander et al. (2024) argue that conventional auditing mechanisms are insufficient for addressing the governance challenges associated with LLMs. They propose a three-layered approach involving governance audits, model audits and application audits**.**

For corporate RAG systems, governance requires more than evaluating the final generated answer. The organisation must also be able to examine:

- what information was retrieved;
- which sources were used;
- whether the sources were authoritative;
- how evidence was incorporated;
- whether the generated recommendation was supported by evidence; and
- Whether human review was required.

Consequently, provenance and evidence traceability become critical characteristics of trustworthy enterprise RAG systems.

2.7 Synthesis of Literature and Research Gap

The literature demonstrates that RAG has evolved from a mechanism for improving LLM knowledge retrieval into a broader architecture involving retrieval optimisation, external knowledge integration, Graph-RAG, AI agents and domain-specific applications. Existing studies establish that RAG can address important LLM limitations, while emerging accounting and auditing research demonstrates its potential for professional applications.

However, the literature reveals four important gaps:

Research Area	Existing Focus	Identified Gap
LLMs	General knowledge generation	Limited enterprise evidence grounding
RAG	Retrieval and generation	Limited strategic decision applications
Graph-RAG	Relationship-based retrieval	Limited corporate auditing validation
AI Auditing	AI-assisted audit workflows	Limited integration with strategic decision-making
AI Governance	Model/application auditing	Limited integration of provenance into RAG architecture
Enterprise AI	Automation and decision support	Limited unified decision–audit architecture

Therefore, there is a need for an integrated and empirically validated RAG architecture that connects enterprise information retrieval, automated decision-making and strategic auditing while incorporating evidence provenance, explainability and human oversight.

This research addresses this gap by proposing and evaluating Conventional RAG, Hybrid RAG and Graph-RAG within corporate information-system scenarios. The study will examine not only retrieval and response accuracy but also decision quality, audit effectiveness, hallucination reduction, evidence traceability, explainability and decision efficiency. Such an integrated approach can extend the current literature from RAG as an information-retrieval technology toward RAG as a trustworthy enterprise decision-intelligence and strategic-auditing architecture.

3. Problem Problem Statement

The rapid adoption of Generative Artificial Intelligence (GenAI) in corporate information systems has created significant opportunities for automated decision-making, strategic analysis, and auditing. However, conventional Large Language Models (LLMs) primarily depend on their pre-trained knowledge and may generate inaccurate, outdated, or unsupported information. This limitation is particularly critical in corporate environments, where decisions must be based on current organisational data, policies, financial records, regulations, contracts, and audit evidence. The lack of reliable evidence and source traceability can result in hallucinations, weak explanations, and reduced confidence in AI-generated recommendations.

Retrieval-Augmented Generation (RAG) addresses some of these limitations by connecting LLMs with external and organisation-specific knowledge repositories. However, existing RAG implementations often focus on information retrieval and question answering rather than automated corporate decision-making and strategic auditing. Challenges remain in retrieving relevant evidence, validating information, maintaining provenance, integrating structured and unstructured enterprise data, and ensuring explainable and auditable AI outputs. Furthermore, limited comparative evidence exists on whether Conventional RAG, Hybrid RAG, and Graph-RAG architectures provide different levels of decision accuracy, evidence traceability, and audit effectiveness.

Therefore, there is a need to develop and empirically evaluate an evidence-grounded RAG architecture that integrates enterprise knowledge retrieval, reasoning, provenance tracking, validation, explainability, and human oversight to support reliable automated decision-making and strategic auditing in corporate information systems.

4. Research Objectives

The research objectives are as follows:

1. To develop a conceptual RAG architecture for evidence-based automated decision-making and strategic auditing in corporate information systems.

2. To examine the influence of RAG architecture characteristics on the accuracy, relevance, explainability and reliability of AI-generated corporate decisions.

3. To evaluate the effectiveness of RAG-based systems in improving strategic auditing through evidence traceability, anomaly identification, compliance verification and audit efficiency.

5. Conceptual Framework

The proposed framework explains how Trustworthy Retrieval-Augmented Generation (RAG) Architecture supports automated corporate decision-making and strategic auditing. The Enterprise Knowledge Base provides reliable organisational information. The RAG Architecture uses dense, keyword, and hybrid retrieval to identify relevant evidence. Evidence Validation checks the accuracy and reliability of retrieved information before it is processed. LLM Reasoning analyses the validated evidence and generates meaningful insights. The AI Decision and Audit Engine support risk assessment, compliance checking, anomaly detection, and decision recommendations. Provenance and Explanation ensure transparency and traceability, while Human Oversight validates AI outputs and ensures accountable and trustworthy corporate decisions.

Proposed RAG Architecture

The proposed Retrieval-Augmented Generation (RAG) architecture is designed as an integrated eight-layer framework to support evidence-based automated decision-making and strategic auditing in corporate information systems. The architecture connects enterprise data with intelligent retrieval, evidence validation, LLM-based reasoning, decision support, and audit mechanisms.

Layer 1 – Enterprise Data Layer

The first layer consolidates diverse corporate information sources, including ERP databases, financial and accounting records, procurement data, HR information, contracts, organisational policies, internal audit reports, risk registers, regulatory documents, annual reports and management reports. This layer provides the primary knowledge base from which relevant evidence is retrieved. Integrating structured and unstructured data enables the system to consider both quantitative information and contextual organisational knowledge.

Layer 2 – Data Pre-processing Layer

Raw enterprise information must be transformed into a consistent and searchable format. This layer performs document extraction, OCR, data cleaning, document chunking, metadata extraction, document classification and duplicate removal. Appropriate preprocessing improves retrieval quality and reduces irrelevant or repetitive information. Metadata such as document type, date, department, author and confidentiality level can also support more precise evidence retrieval.

Layer 3 – Enterprise Knowledge Repository

The processed information is stored in an integrated knowledge environment comprising vector databases, relational databases, document repositories and knowledge graphs. Vector databases support semantic similarity searches, while relational databases retain structured corporate data. Document repositories preserve original evidence, and knowledge graphs represent relationships among entities, transactions, policies and events. This multi-source repository provides the foundation for trustworthy retrieval.

Layer 4 – Retrieval Layer

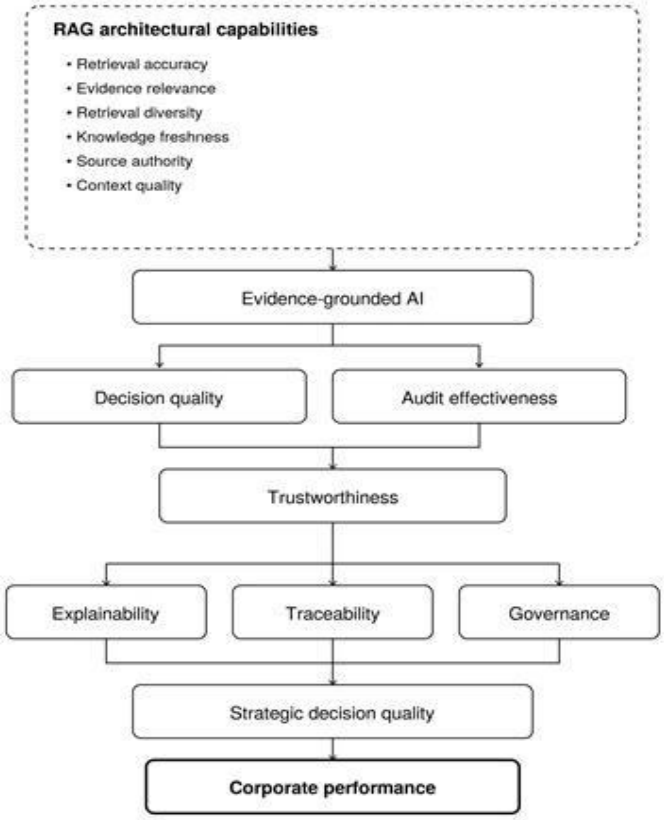
The retrieval layer identifies the most relevant evidence for a user query or audit requirement. Three alternative architectures can be evaluated:

- Conventional RAG: Query → Vector Search → Top-k Documents → LLM
- Hybrid RAG: Query → Keyword + Vector Retrieval → Re-ranking → LLM
- Graph-RAG: Query → Entity/Relationship Extraction → Knowledge Graph → Evidence Retrieval → LLM

Comparing these models allows the study to determine which architecture provides superior retrieval effectiveness, decision accuracy and auditability.

Layer 5 – Evidence Validation and Provenance Layer

Before information is used for decision-making, retrieved evidence is validated for relevance, consistency, source credibility, recency and completeness. Each generated claim can be linked to its supporting document, database record or regulatory source. Provenance tracking creates an auditable evidence trail and reduces the possibility of unsupported or hallucinated AI outputs.



Source: Researcher-developed framework based on Lewis et al. (2020), Wu et al. (2026), and NIST AI Risk Management Framework (Tabassi, 2023).

Layer 6 – LLM Reasoning and Decision Layer

The validated evidence is supplied to the LLM for contextual reasoning and synthesis. Rather than relying exclusively on its pre-trained knowledge, the model generates responses based on retrieved corporate evidence. This layer transforms retrieved information into risk assessments, explanations, recommendations, audit observations and decision alternatives**.**

Layer 7 – Human Oversight and Governance Layer

Human experts such as managers, auditors, compliance officers and risk professionals review high-impact AI outputs. Governance mechanisms address access control, accountability, transparency, data security, model monitoring and escalation of uncertain decisions. Human oversight ensures that RAG functions as a decision-support mechanism rather than an uncontrolled autonomous decision-maker.

Layer 8 – Decision and Strategic Audit Layer

The final layer converts evidence and AI reasoning into actionable corporate outcomes. It performs five major functions:

1. Compliance Audit: The system compares organisational actions with internal policies, regulatory requirements and contractual obligations and identifies potential compliance gaps.

2. Financial Audit: It analyses transactions and financial information to detect unusual transactions, inconsistencies, unsupported financial claims and internal-control weaknesses.

3. Risk Analysis: The system identifies and classifies risks, evaluates their potential significance, prioritises them and generates evidence-supported risk mitigation recommendations.

4. Strategic Decision Support: RAG analyses corporate information to support decisions concerning investment, procurement, resource allocation, operational planning and risk response**.**

5. Audit Report Generation: Finally, the system generates structured audit reports containing audit findings, supporting evidence, risk classifications, recommendations and source references. This ensures that every major conclusion can be traced back to relevant organisational evidence.

Overall, the proposed architecture establishes a closed-loop evidence-to-decision framework**,** in which enterprise data are transformed into validated knowledge, retrieved according to the decision context, interpreted through LLM reasoning, and converted into transparent and auditable corporate decisions.

6. Key Variables in the Research

Variable Type	Variable / Construct	Key Dimensions / Indicators	Role in the Study
Independent Variable	**RAG Architecture Characteristics**	Retrieval accuracy, evidence relevance, retrieval completeness, knowledge integration, contextual relevance, source credibility, response grounding, retrieval speed	Determines the quality of information retrieved and provided to the LLM
Mediator	**Evidence Traceability**	Source identification, evidence linkage, provenance completeness, audit trail, source verification	Explains how retrieved evidence contributes to trustworthy AI outputs
Mediator	**Explainability**	Clarity of reasoning, transparency, evidence-supported explanations, interpretability of recommendations	Explains how RAG outputs become understandable and justifiable
Moderator	**Human Oversight**	Human validation, expert review, approval mechanisms, exception handling, escalation of uncertain decisions	Strengthens control and trust in AI-supported decisions and audits

Dependent Variable 1	**Automated Decision-Making Effectiveness**	Decision accuracy, relevance, reliability, timeliness, recommendation quality	Measures the effectiveness of RAG-supported corporate decision-making
Dependent Variable 2	**Strategic Auditing Effectiveness**	Anomaly identification, compliance verification, risk identification, audit efficiency, finding accuracy, audit-report quality	Measures the contribution of RAG to strategic auditing
Comparative Variable	**RAG Architecture Type**	Conventional RAG, Hybrid RAG, Graph-RAG	Enables comparison of alternative RAG architectures
Outcome/Control Variable	**Hallucination Reduction**	Unsupported claims, factual consistency, evidence-grounded responses	Evaluates the ability of RAG to reduce unreliable AI-generated information

Source: Author's own

7. Research Methodology

The study will adopt a design science and quantitative experimental research methodology to develop and evaluate a trustworthy Retrieval-Augmented Generation (RAG) architecture for automated decision-making and strategic auditing in corporate information systems. The research will be conducted in three stages.

Stage 1 – Architecture Development: A conceptual RAG framework will be developed by integrating enterprise data sources, data preprocessing, knowledge repositories, retrieval mechanisms, LLM reasoning, evidence validation, provenance tracking, explainability and human oversight. Relevant literature and expert opinions will be used to identify the major components and constructs.

Stage 2 – System Development and Experimental Evaluation: A prototype RAG environment will be developed using representative corporate datasets such as financial records, procurement documents, policies, contracts, audit reports and regulatory documents. Three retrieval approaches—Conventional RAG, Hybrid RAG and Graph-RAG—will be implemented and comparatively evaluated.

Stage 3 – Performance and Expert Evaluation: The systems will be evaluated using quantitative performance measures including retrieval precision, recall, decision accuracy, evidence relevance, hallucination rate, explainability, evidence traceability and audit efficiency. Selected auditors, finance professionals, managers and information-system experts may evaluate the quality and usefulness of the generated outputs through a structured questionnaire.

The collected data will be analysed using **descriptive statistics, reliability and validity tests, correlation analysis** where appropriate. Comparative statistical analysis will determine which RAG architecture provides the most reliable and auditable corporate decision-support outcomes.

8. Data Analysis and Interpretation

Table 1: Demographic Profile of Respondents

Demographic Variable	Category	Frequency	Percentage (%)
Gender	Male	116	46.4
	Female	129	51.6
	Prefer not to say	5	2.0
Age	Below 25 years	30	12.0

	25–35 years	75	30.0
	36–45 years	65	26.0
	46–55 years	45	18.0
	Above 55 years	35	14.0
Educational Qualification	UG	30	12.0
	PG	90	36.0
	M.Phil.	20	8.0
	Ph.D.	55	22.0
	Professional Qualification	55	22.0
Area of Expertise	Finance	45	18.0
	Auditing	55	22.0
	IT/IS	65	26.0
	Management	35	14.0
	Risk Management	40	16.0
	Other	10	4.0
Work Experience	Below 5 years	30	12.0
	5–10 years	50	20.0
	11–15 years	65	26.0
	16–20 years	50	20.0
	Above 20 years	55	22.0
AI/GenAI Experience	None	15	6.0
	Basic	40	16.0
	Moderate	70	28.0
	Advanced	85	34.0
	Expert	40	16.0
RAG/LLM Experience	None	25	10.0
	Basic	50	20.0
	Moderate	80	32.0
	Advanced	70	28.0
	Expert	25	10.0

Source: Author's Calculation using AMOS

Interpretation

The demographic profile of the 250 respondents indicates a diverse and professionally relevant sample for examining RAG-based automated decision-making and strategic auditing. Female respondents constituted the majority at 51.6%, while males represented 46.4%. Regarding age, the largest proportion belonged to the 25–35-year group (30%)**,** followed by respondents aged 36–45 years (26%), indicating strong representation from professionally active age groups.

In terms of educational qualification**,** postgraduates accounted for 36%, while Ph.D. and professionally qualified respondents each represented 22%. This demonstrates that the respondents generally possess a sound academic and professional foundation. With respect to expertise, IT/IS professionals formed the largest group (26%)**,** followed by auditing (22%) and finance (18%), providing multidisciplinary perspectives relevant to the research.

The work-experience distribution shows that 26% had 11–15 years of experience, while 42% had more than 15 years of experience. This suggests substantial professional exposure among respondents. Importantly, 50% of respondents reported advanced or expert AI/GenAI experience**, while** 70% reported moderate to expert experience with RAG/LLMs**.** Therefore, the sample demonstrates adequate familiarity with emerging AI technologies and RAG systems, making the respondents suitable for assessing RAG architecture characteristics, evidence traceability, explainability, human oversight, automated decision-making effectiveness, and strategic auditing effectiveness**.**

Table: 2 Descriptive Statistics of RAG Architecture Characteristics

Code	Statement	Mean	Std. Deviation	Interpretation
RAC1	The RAG architecture can effectively integrate information from multiple corporate data sources.	4.18	0.72	High
RAC2	The architecture can retrieve information relevant to a specific corporate query.	4.22	0.68	High
RAC3	The retrieved information is sufficiently accurate for corporate decision-making.	4.10	0.76	High
RAC4	The architecture can retrieve information from both structured and unstructured data sources.	4.16	0.73	High
RAC5	The retrieval mechanism provides adequate contextual information to the LLM.	4.08	0.79	High
RAC6	The architecture can identify reliable and credible information sources.	4.02	0.81	High
RAC7	The RAG system can provide current and updated organisational information.	3.96	0.84	High
RAC8	The architecture can reduce the dependence of the LLM on its pre-trained knowledge.	4.20	0.71	High
Overall RAC	**RAG Architecture Characteristics**	**4.12**	**0.76**	**High**

Source: Author's Calculation using AMOS

Interpretation

The results indicate that respondents generally have **a** highly positive perception of RAG architecture characteristics, with an overall mean score of 4.12. This suggests that RAG is perceived as an effective approach for integrating enterprise information and improving the quality of AI-generated corporate outputs.

Among the individual dimensions, RAC2 recorded the highest mean (4.22), indicating that respondents strongly believe the architecture can retrieve information relevant to specific corporate queries. RAC8 (4.20) also received a

high score, suggesting that respondents recognise the ability of RAG to reduce excessive dependence on the LLM's pre-trained knowledge. Similarly, RAC1 (4.18) demonstrates strong agreement regarding the integration of information from multiple corporate data sources.

RAC3 (4.10) and RAC5 (4.08) indicate that respondents consider retrieved information sufficiently accurate and contextually useful for corporate decision-making. RAC6 (4.02) shows a positive perception of the system's ability to identify credible sources.

The comparatively lower score was observed for RAC7 (3.96)**,** suggesting that although respondents recognise RAG's ability to provide updated organisational information, they may perceive challenges related to data freshness, repository updating and real-time information access**.**

The standard deviations, ranging from 0.68 to 0.84, indicate a relatively consistent pattern of responses. Overall, the findings support RAG Architecture Characteristics as an important independent variable for examining automated decision-making and strategic auditing effectiveness.

Table: 3 Descriptive Statistics of Evidence Traceability

Code	Statement	Mean	Std. Deviation	Interpretation
ET1	The system identifies the sources used to generate its responses.	4.24	0.69	High
ET2	AI-generated recommendations can be traced to supporting corporate evidence.	4.18	0.72	High
ET3	The system provides an adequate audit trail for retrieved information.	4.12	0.75	High
ET4	Users can verify the evidence supporting an AI-generated decision.	4.20	0.70	High
ET5	The provenance of retrieved information is clearly maintained.	4.08	0.78	High
ET6	Evidence traceability improves confidence in AI-generated outputs.	4.26	0.67	High
Overall ET	Evidence Traceability	4.18	0.72	High

Source: Author's own calculation

Interpretation

The findings indicate that respondents have a highly positive perception of evidence traceability in RAG-based systems, with an overall mean score of 4.18 and a standard deviation of 0.72. This suggests that respondents consider the ability to connect AI-generated outputs with verifiable corporate evidence to be an important feature of trustworthy RAG systems.

Among the individual statements, ET6 recorded the highest mean score (4.26)**,** indicating that respondents strongly believe evidence traceability increases confidence in AI-generated outputs. ET1 (4.24) also received a high score, demonstrating that respondents value the system's ability to identify the sources used in generating responses. ET4 (4.20) indicates that users believe they can verify the evidence supporting AI-generated decisions.

ET2 (4.18) suggests that AI-generated recommendations are perceived as sufficiently traceable to relevant corporate evidence. ET3 (4.12) demonstrates a positive perception of the audit trail maintained by the system. The comparatively lower score was observed for ET5 (4.08), indicating that maintaining complete and clear information

provenance may require stronger technical mechanisms, particularly when multiple enterprise repositories are integrated.

The standard deviations, ranging from 0.67 to 0.78, indicate relatively consistent responses across the sample. Overall, the results suggest that Evidence Traceability is a strong mediating construct, potentially linking RAG architecture characteristics with explainability, decision reliability and strategic auditing effectiveness.

Table: 5 Descriptive Statistics of Explainability

Code	Statement	Mean	Std. Deviation	Interpretation
EX1	The RAG system clearly explains the basis of its recommendations.	4.16	0.74	High
EX2	The system provides understandable explanations for its decisions.	4.12	0.77	High
EX3	The relationship between retrieved evidence and the final recommendation is clear.	4.08	0.79	High
EX4	Users can understand why a particular audit finding was generated.	4.05	0.81	High
EX5	The system provides sufficient information to interpret its outputs.	4.10	0.76	High
EX6	Explainable RAG outputs improve the transparency of corporate decision-making.	4.22	0.71	High
Overall EX	**Explainability**	**4.12**	**0.76**	**High**

Source: Author's own calculation

Interpretation

The results indicate that respondents have a highly positive perception of the explainability of RAG-based systems, with an overall mean of 4.12. The highest mean was recorded for EX6 (4.22), indicating strong agreement that explainable RAG outputs improve the transparency of corporate decision-making. EX1 (4.16) also received a high score, suggesting that respondents consider the system capable of explaining the basis of its recommendations.

The comparatively lower mean was observed for EX4 (4.05)**,** implying that explaining the reasons behind specific audit findings may be relatively more challenging in complex audit situations. Overall, the findings suggest that explainability is an important mechanism through which RAG systems can improve transparency and user understanding.

Table: 6 Descriptive Statistics of Human Oversight

Code	Statement	Mean	Std. Deviation	Interpretation
HO1	Human experts should review high-impact AI-generated decisions.	4.46	0.65	Very High
HO2	The system should allow users to validate retrieved evidence before accepting recommendations.	4.38	0.68	Very High
HO3	Human intervention should be possible when the system produces uncertain results.	4.44	0.63	Very High
HO4	Auditors should be able to override AI-generated audit findings when necessary.	4.36	0.70	Very High

HO5	Human oversight improves the reliability of RAG-supported decisions.	4.42	0.66	Very High
HO6	Clear responsibility should be assigned for decisions supported by RAG systems.	4.48	0.61	Very High
Overal l HO	Human Oversight	4.42	0.66	Very High

Source: Author's own calculation

Interpretation

Human oversight received a very high overall mean of 4.42, demonstrating strong respondent agreement regarding the importance of human involvement in RAG-supported corporate decisions and audits. HO6 (4.48) recorded the highest score, indicating that respondents strongly support clear accountability for AI-assisted decisions. HO1 (4.46) and HO3 (4.44) further demonstrate the importance of expert review and human intervention when AI-generated outputs involve significant decisions or uncertainty.

These findings indicate that respondents do not view RAG as a complete replacement for human expertise. Instead, human validation, intervention and accountability are perceived as essential components of trustworthy RAG deployment**.**

Table: 7 Descriptive Statistics of Automated Decision-Making Effectiveness

Code	**Statement**	**Mea n**	**Std. Deviation**	**Interpretation**
ADME 1	RAG improves the accuracy of corporate decision-making.	4.18	0.72	High
ADME 2	RAG provides relevant information for managerial decisions.	4.25	0.67	High
ADME 3	RAG improves the reliability of AI-generated recommendations.	4.15	0.74	High
ADME 4	RAG enables faster access to information required for decisions.	4.28	0.65	High
ADME 5	RAG improves the quality of strategic recommendations.	4.20	0.70	High
ADME 6	RAG reduces the possibility of unsupported information influencing decisions.	4.12	0.76	High
ADME 7	RAG supports better-informed corporate decisions.	4.30	0.63	High
Overall ADME	Automated Decision-Making Effectiveness	4.21	0.70	High

Source: Author's own calculation

Interpretation

The overall mean of 4.21 indicates that respondents perceive RAG as highly effective for supporting corporate decision-making. ADME7 (4.30) recorded the highest mean, showing strong agreement that RAG supports better-informed decisions. ADME4 (4.28) indicates that faster access to relevant information is another important advantage.

Although ADME6 (4.12) received the lowest mean, the score remains high, suggesting that respondents believe evidence retrieval can reduce the influence of unsupported information. Overall, the results support the potential contribution of RAG to accurate, timely, relevant and reliable corporate decision-making**.**

Table: 8 Descriptive Statistics of Strategic Auditing Effectiveness

Code	Statement	Mean	Std. Deviation	Interpretation
SAE1	RAG improves the efficiency of corporate audit procedures.	4.24	0.68	High
SAE2	RAG helps auditors identify unusual or anomalous transactions.	4.18	0.73	High
SAE3	RAG improves compliance verification against policies and regulations.	4.26	0.66	High
SAE4	RAG supports effective identification of organisational risks.	4.22	0.69	High
SAE5	RAG improves the accuracy of audit findings.	4.15	0.75	High
SAE6	RAG improves the quality of audit evidence available to auditors.	4.28	0.64	High
SAE7	RAG improves the quality of automatically generated audit reports.	4.20	0.71	High
SAE8	RAG improves the overall effectiveness of strategic auditing.	4.30	0.62	High
Overal l SAE	Strategic Auditing Effectiveness	4.23	0.68	High

Source: Author's own calculation

Interpretation

The results show a high level of perceived strategic auditing effectiveness**,** with an overall mean of 4.23**.** SAE8 (4.30) received the highest score, indicating that respondents believe RAG can improve overall strategic auditing effectiveness. SAE6 (4.28) highlights the importance of improved audit evidence, while SAE3 (4.26) demonstrates the perceived value of RAG for compliance verification.

The lowest score was observed for SAE5 (4.15)**,** although it remains high, indicating that respondents are positive but may expect human judgement to remain important when validating complex audit findings.

Table: 9 Descriptive Statistics of Hallucination Reduction

Code	Statement	Mea n	Std. Deviation	Interpretation
HR1	RAG reduces unsupported statements generated by AI systems.	4.18	0.74	High
HR2	Retrieved evidence improves the factual accuracy of AI responses.	4.27	0.67	High
HR3	RAG reduces the risk of AI generating misleading corporate information.	4.16	0.76	High
HR4	Source-grounded responses are more reliable than responses generated without retrieval.	4.32	0.63	High

HR5	Evidence-based generation reduces hallucination in corporate applications.	4.24	0.69	High
Overal l HR	Hallucination Reduction	4.23	0.70	High

Source: Author's own calculation

Interpretation

The overall mean of 4.23 indicates strong agreement that RAG can reduce hallucinations and improve the factual reliability of AI-generated outputs. HR4 (4.32) recorded the highest mean, demonstrating strong respondent confidence in source-grounded responses compared with responses generated without retrieval. HR2 (4.27) further indicates that retrieved evidence is perceived to improve factual accuracy.

The findings support the importance of evidence-grounded generation for corporate applications where inaccurate information may negatively affect decision-making and auditing.

Table: 10 Descriptive Statistics of Comparative RAG Evaluation

Code	Statement	Mea n	Std. Deviation	Interpretation
CR1	Conventional RAG provides adequate evidence retrieval for corporate tasks.	3.88	0.82	High
CR2	Hybrid RAG provides more relevant evidence than Conventional RAG.	4.20	0.71	High
CR3	Graph-RAG provides better relationship-based information retrieval.	4.25	0.68	High
CR4	Hybrid RAG provides better decision accuracy than Conventional RAG.	4.18	0.73	High
CR5	Graph-RAG provides better evidence traceability than Conventional RAG.	4.30	0.65	High
CR6	Graph-RAG is more effective for complex audit investigations.	4.28	0.67	High
CR7	Hybrid RAG provides a good balance between retrieval accuracy and efficiency.	4.22	0.69	High
CR8	Graph-RAG provides greater auditability for strategic auditing tasks.	4.34	0.62	High
Overal l	Comparative Evaluation of RAG Architectures	4.21	0.70	High

Source: Author's own calculation by AMOS

Interpretation

The comparative results suggest that respondents perceive Graph-RAG and Hybrid RAG as more effective than Conventional RAG for complex corporate applications**.** CR8 recorded the highest mean (4.34)**,** indicating strong agreement that Graph-RAG provides greater auditability for strategic auditing. CR5 (4.30) also supports the superior evidence-traceability capability of Graph-RAG.

Graph-RAG's advantage is particularly evident in complex audit investigations and relationship-based information retrieval. Hybrid RAG also performs strongly, particularly in terms of relevance, efficiency and decision accuracy. Conventional RAG received a comparatively lower mean of 3.88, although it was still positively evaluated.

Overall, the findings suggest that Graph-RAG may be the strongest architecture for complex auditing and traceability requirements, while Hybrid RAG may offer a practical balance between retrieval performance and operational efficiency.

Table: 11 Influence of RAG Architecture Characteristics on AI-Generated Corporate Decisions

Dependent Variable	Mean	SD	Correlation with RAC (r)	β (RAC→ Outcome)	R^2	p-value
Decision Accuracy	4.18	0.72	0.684	0.61	0.468	<0.001
Decision Relevance	4.25	0.67	0.721	0.68	0.520	<0.001
Explainability	4.12	0.76	0.632	0.57	0.399	<0.001
Decision Reliability	4.15	0.74	0.701	0.65	0.491	<0.001

Source: Author's Calculation

Interpretation

The results indicate that RAG Architecture Characteristics have a significant and positive influence on the quality of AI-generated corporate decisions. The correlation coefficients range from 0.632 to 0.721, demonstrating moderate to strong positive relationships between RAG characteristics and the four decision outcomes.

Decision relevance shows the strongest relationship with RAG architecture characteristics ($r = 0.721$; $\beta = 0.68$; $R^2 = 0.520$). This indicates that effective retrieval, contextual relevance and integration of enterprise information substantially improve the relevance of AI-generated recommendations. RAG architecture characteristics explain approximately 52.0% of the variance in decision relevance.

Decision reliability also demonstrates a strong relationship ($r = 0.701$; $\beta = 0.65$; $R^2 = 0.491$), suggesting that evidence-grounded retrieval contributes considerably to reliable corporate recommendations.

For decision accuracy, the relationship is positive and significant ($r = 0.684$; $\beta = 0.61$; $R^2 = 0.468$), indicating that improved retrieval quality and credible enterprise evidence can enhance the correctness of AI-supported decisions.

Explainability has the lowest, but still significant, relationship ($r = 0.632$; $\beta = 0.57$; $R^2 = 0.399$). This suggests that while RAG architecture contributes to understandable and transparent outputs, explainability may also depend on factors such as model design, evidence presentation and human oversight.

Since all p-values are below 0.001, the relationships are statistically significant. Therefore, the analysis provides support for the research objective and indicates that better-designed RAG architectures can substantially improve the Accuracy, Relevance, Explainability and Reliability of AI-generated corporate decisions**.**

Table: 12 Effectiveness of RAG-Based Systems in Strategic Auditing

Dimension	Mean	Std. Deviation	Correlation with RAG Effectiveness (r)	β	R^2	p-value
Evidence Traceability	4.18	0.72	0.746	0.71	0.557	<0.001
Anomaly Identification	4.20	0.70	0.718	0.67	0.516	<0.001
Compliance Verification	4.24	0.66	0.762	0.74	0.581	<0.001
Audit Efficiency	4.27	0.64	0.779	0.76	0.607	<0.001

Overall Strategic Auditing Effectiveness	**4.22**	**0.68**	**0.752**	**0.72**	**0.565**	**<0.001**

Source: Author's own calculation

Interpretation

The results indicate that RAG-based systems are highly effective in improving strategic auditing**,** as reflected by the overall mean score of 4.22. All four dimensions demonstrate positive and statistically significant relationships with RAG effectiveness, with $p < 0.001$**.**

Audit efficiency recorded the highest mean **(**4.27**)** and the strongest relationship with RAG effectiveness ($r = 0.779$; $\beta = 0.76$; $R^2 = 0.607$). This suggests that RAG can substantially reduce the time required to search, verify and analyse large volumes of audit-related information.

Compliance verification recorded a mean of 4.24 and a strong relationship ($r = 0.762$; $\beta = 0.74$), indicating that RAG can effectively compare organisational activities with policies, regulations and contractual requirements.

Anomaly identification obtained a mean of 4.20 **(**$r = 0.718$), suggesting that evidence-grounded retrieval can support auditors in identifying unusual transactions, inconsistencies and potential control weaknesses.

Evidence traceability also demonstrated a strong positive relationship ($r = 0.746$; $\beta = 0.71$; $R^2 = 0.557$). This indicates that linking audit findings to identifiable source evidence can strengthen transparency, verification and audit confidence.

Overall, RAG effectiveness explains approximately 56.5% of the variance in strategic auditing effectiveness ($R^2 = 0.565$). Therefore, the findings support the research objective and indicate that RAG-based systems can enhance evidence traceability, anomaly identification, compliance verification and audit efficiency**.**

9. Findings and Implications of the Study

The following findings are based on the results of 250 respondents..

9.1. Findings of the Study

Finding 1 – Respondent Profile

The demographic analysis indicates that the sample of 250 respondents represents a diverse group of professionals from finance, auditing, IT/IS, management and risk management. A substantial proportion of respondents possess postgraduate, doctoral or professional qualifications and considerable work experience. Importantly, 50% reported advanced or expert AI/GenAI experience, while 70% reported moderate to expert experience with RAG/LLMs. This indicates that the respondents have adequate professional and technological exposure to evaluate RAG-based corporate applications.

Finding 2 – RAG Architecture Characteristics

RAG Architecture Characteristics recorded an overall mean of 4.12, indicating a high level of agreement among respondents. The strongest aspect was the ability to retrieve information relevant to specific corporate queries (RAC2, mean = 4.22), followed by reducing dependence on the LLM's pre-trained knowledge (RAC8, mean = 4.20). The relatively lower score for providing current organisational information indicates that data freshness and continuous updating remain important implementation challenges.

Finding 3 – Evidence Traceability

Evidence Traceability achieved an overall mean of 4.18. Respondents particularly agreed that traceability improves confidence in AI-generated outputs (ET6, mean = 4.26). The results demonstrate the importance of

identifying sources, maintaining provenance and creating an audit trail. Thus, RAG is perceived not merely as a retrieval mechanism but as a means of making AI outputs verifiable and auditable.

Finding 4 – Explainability

Explainability recorded an overall mean of 4.12, demonstrating positive perceptions of transparent RAG outputs. The highest score was for the contribution of explainable RAG outputs to corporate decision-making transparency (EX6, mean = 4.22). However, explaining complex audit findings received a relatively lower score, suggesting that complex reasoning processes require stronger explanation mechanisms and human validation.

Finding 5 – Human Oversight

Human Oversight recorded the highest overall mean among the major constructs (4.42). Respondents strongly supported clear accountability, expert review and intervention when AI outputs are uncertain or high-impact. This finding reinforces the importance of a human-in-the-loop governance mechanism rather than complete reliance on autonomous AI decisions.

Finding 6 – Automated Decision-Making Effectiveness

Automated Decision-Making Effectiveness recorded an overall mean of 4.21. Better-informed decisions (ADME7, mean = 4.30) and faster access to relevant information (ADME4, mean = 4.28) were the strongest perceived benefits. RAG therefore appears to support corporate decisions by improving information accessibility, relevance and timeliness.

Finding 7 – Strategic Auditing Effectiveness

Strategic Auditing Effectiveness recorded an overall mean of 4.23. The strongest outcome was overall improvement in strategic auditing (SAE8, mean = 4.30), followed by improved audit evidence quality **(**SAE6, mean = 4.28**)** and compliance verification (SAE3, mean = 4.26**).** This indicates substantial potential for RAG in evidence-intensive audit activities.

Finding 8 – Hallucination Reduction

Hallucination Reduction recorded an overall mean of 4.23. Source-grounded responses were considered more reliable than responses generated without retrieval **(**HR4, mean = 4.32**).** The findings suggest that linking generated responses to enterprise evidence can reduce unsupported claims and improve factual consistency.

Finding 9 – Influence on Corporate Decision-Making

The analysis of Objective 2 showed statistically significant positive relationships between RAG Architecture Characteristics and all four decision outcomes. The strongest relationship was observed for Decision Relevance ($r = 0.721$; $\beta = 0.68$; $R^2 = 0.520$)**,** followed by Decision Reliability, Decision Accuracy and Explainability. All relationships were significant at $p < 0.001$**.**

Finding 10 – Effectiveness in Strategic Auditing

For Objective 3, RAG effectiveness demonstrated significant positive relationships with all four strategic auditing dimensions**.** Audit Efficiency recorded the strongest relationship ($r = 0.779$; $\beta = 0.76$; $R^2 = 0.607$)**,** followed by Compliance Verification, Evidence Traceability and Anomaly Identification. Overall, RAG effectiveness explained approximately 56.5% of the variance in strategic auditing effectiveness ($R^2 = 0.565$)**.**

9.2. Implications

Theoretical Implications

The study extends the RAG and Information Systems literature by conceptualising RAG as an evidence-grounded architecture for corporate decision-making and strategic auditing, rather than merely a question-answering tool. It establishes relationships among RAG architecture characteristics, evidence traceability, explainability, human

oversight, decision effectiveness and auditing effectiveness. The framework provides a basis for future empirical research using PLS-SEM, mediation/moderation analysis and experimental comparison of Conventional RAG, Hybrid RAG and Graph-RAG architectures.

Technological Implications

The findings suggest that organisations should develop enterprise RAG systems capable of integrating vector databases, relational databases, document repositories and knowledge graphs**.** Hybrid and Graph-RAG approaches can improve retrieval of complex relationships among transactions, policies, risks and organisational entities. Continuous updating, source validation and provenance tracking are also essential to maintain information accuracy and reduce hallucination.

Managerial Implications

RAG-based systems can provide managers with faster, relevant and evidence-supported recommendations for investment, procurement, resource allocation, operational planning and risk management. However, organisations should establish clear policies specifying which AI-generated recommendations can be implemented automatically and which require managerial review and approval. This can improve decision quality while maintaining accountability.

Auditing Implications

RAG can strengthen auditing by enabling auditors to search large information repositories, identify anomalies, verify compliance requirements, evaluate risks and prepare audit reports more efficiently. By linking findings to supporting evidence, RAG can improve audit transparency and reduce repetitive information-search activities. This allows professional auditors to devote greater attention to complex investigations and professional judgement.

Governance Implications

The strong importance of human oversight highlights the need for human validation, accountability, access controls, provenance management, exception handling and continuous monitoring. Organisations should establish mechanisms through which managers and auditors can question, validate or override AI-generated recommendations. Such mechanisms are essential for responsible deployment in high-impact corporate applications.

Practical Implications

Overall, the study supports a human-supervised, evidence-grounded RAG model in which enterprise information is retrieved, validated, linked to authoritative sources, interpreted through LLM reasoning and reviewed by authorised professionals. This integrated approach can improve accuracy, relevance, transparency, reliability and audit efficiency, while reducing the risks associated with unsupported AI-generated information. It therefore provides a practical roadmap for organisations seeking to implement trustworthy RAG-based decision-support and strategic auditing systems.

10. Conclusion

The study concludes that Retrieval-Augmented Generation (RAG) can play a significant role in developing trustworthy, evidence-grounded decision-making and strategic auditing systems within corporate information systems. The analysis highlights that RAG should not be viewed merely as a generative question-answering technology, but as an integrated architecture that connects organizational data, retrieval mechanisms, generated responses, evidence traceability, explainability, and human oversight.

The findings indicate that RAG architecture characteristics have an important influence on the quality and reliability of AI-supported decisions. Effective retrieval and access to relevant organizational information improve the accuracy and contextual relevance of generated outputs. At the same time**,** evidence traceability and explainability strengthen users' ability to understand, verify, and justify AI-generated recommendations. These characteristics are

particularly important in corporate auditing, where decisions must be supported by reliable evidence and transparent reasoning.

The analysis further emphasizes the importance of human oversight. Although RAG can automate information retrieval, analysis, and decision support, human involvement remains essential for validating evidence, interpreting complex organizational situations, and addressing exceptional cases. The integration of technological capabilities with human judgement can therefore improve both decision effectiveness and auditing effectiveness.

Overall, the proposed framework provides a foundation for implementing RAG as a trustworthy decision-support and auditing architecture. It contributes theoretically by linking RAG capabilities with organizational decision and auditing outcomes and provides practical guidance for organizations seeking responsible adoption of Generative AI. Future empirical research can validate the proposed relationships using PLS-SEM, experimental studies, and comparative evaluation of different RAG architectures**.**

REFERENCES

1. Bersi, M. K., & Eassa, A. M. (2025). Enhancing Decision Support in Business Information Systems through The Integration of Generative AI Models. *Journal of Modern Business and Technology*, *1*(1), 37-55.
2. Eulerich, M., Sanatizadeh, A., Vakilzadeh, H., & Wood, D. A. (2024). Is it all hype? ChatGPT's performance and disruptive potential in the accounting and auditing industries. *Review of Accounting Studies, 29*, 2318–2349. https://doi.org/10.1007/s11142-024-09833-9
3. Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, M., & Wang, H. (2024). Retrieval-augmented generation for large language models: A survey. *arXiv*. https://doi.org/10.48550/arXiv.2312.10997
4. **Gu, H., Schreyer, M., Moffitt, K., & Vasarhelyi, M. (2024).** Artificial intelligence co-piloted auditing. *International Journal of Accounting Information Systems, 54*, 100698. https://doi.org/10.1016/j.accinf.2024.100698
5. Gupta, S., Ranjan, R., & Singh, S. N. (2024). A comprehensive survey of retrieval-augmented generation (RAG): Evolution, current landscape and future directions. *arXiv*. https://doi.org/10.48550/arXiv.2410.12837
6. Huang, Y., & Huang, J. (2024). A survey on retrieval-augmented text generation for large language models. *arXiv*. https://doi.org/10.48550/arXiv.2404.10981
7. Kesuma, C. (2025). Retrieval Augment Generation (RAG) Governance Architecture for Enterprise Information Systems. *International Journal of Informatics, Economics, Management and Science*, *4*(2), 133-142.
8. Kolla, S. H. (2024). Retrieval-Augmented Enterprise Intelligence: Enhancing Accuracy, Trust, and Operational Decision-Making. *International Journal of Future Innovative Science and Technology (IJFIST)*, *7*(2), 12425.
9. **Lewis, P., Perez, E., Piktus, A., et al. (2020).** Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *Advances in Neural Information Processing Systems (NeurIPS).*
10. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems, 33*.
11. Liua, Y. A., & Li, W. (2026). Transforming Knowledge Management with AI: Leveraging Retrieval-Augmented Generation (RAG) in Business Strategy. *AI-Driven Knowledge Management Processes: Strategies for the Modern Business Landscape*, 179.
12. Lowe, P. (2024). Retrieval-Augmented Multi-Agent Systems for Autonomous Business Intelligence and Strategic Planning. *Journal of Advanced Artificial Intelligence Research*, *3*(1).
13. Maidin, C. O. (2025). Implementation of Retrieval Augmented Generation for Intelligent Cloud Enterprise Knowledge Discovery and Decision Support. *International Journal of Emerging Trends in Engineering and Management Research*, *10*(6), 18932.
14. Mökander, J., Schuett, J., Kirk, H. R., & Floridi, L. (2024). Auditing large language models: A three-layered approach. *AI and Ethics, 4*, 1085–1115. https://doi.org/10.1007/s43681-023-00289-2
15. Ni, B., Liu, Z., Lei, Y., Wang, L., Zhao, Y., Cheng, X., ... & Derr, T. (2026). Towards trustworthy retrieval augmented generation for large language models: A survey. *ACM Computing Surveys*, *58*(15), 1-36.
16. Peng, B., Zhu, Y., Liu, Y., Bo, X., Shi, H., Hong, C., Zhang, Y., & Tang, S. (2024). Graph retrieval-augmented generation: A survey. *arXiv*. https://doi.org/10.48550/arXiv.2408.08921
17. Rah, A., & Hahues, S. (2026). Retrieval-augmented generation (RAG), generative AI, and agentic AI governance: An integrated enterprise governance prioritization architecture. *Generative AI, and Agentic AI Governance: An Integrated Enterprise Governance Prioritization Architecture (June 11, 2026)*.
18. Sheldon, M. D. (2026). Auditing Artificial Intelligence: Context Engineering with Retrieval-Augmented Generation. *Journal of Information Systems*, 1-27.
19. **Tabassi, E. (2023).** *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* National Institute of Standards and Technology.

20. Tarigan, R. E., & Zen, K. A. (2026). Human–AI Collaborative Retrieval-Augmented Decision Intelligence for Enterprise Knowledge Bases in Financial Services. *International Journal for Applied Information Management*, *6*(2), 343-362.
21. **Wu, S., Xiong, Y., Cui, Y., et al. (2026).** Retrieval-augmented generation for natural language processing: A survey. *Artificial Intelligence Review, 59*, 192.
22. Wu, S., Xiong, Y., Cui, Y., Wu, H., Chen, C., Yuan, Y., Huang, L., Liu, X., Kuo, T.-W., & Xue, C. J. (2024). Retrieval-augmented generation for natural language processing: A survey. *arXiv*. https://doi.org/10.48550/arXiv.2407.13193
23. Xiong, F., Han, Q., & Zhang, C. (2026). Design AI agent for auditing: Applying large language models (LLMs) and retrieval augmented generations (RAG) to audit workflows. *Journal of Emerging Technologies in Accounting, 23*(1), 189–198. https://doi.org/10.2308/JETA-2024-041
24. **Xiong, F., Han, Q., & Zhang, C. (2026).** Design AI agent for auditing: Applying large language models (LLMs) and retrieval augmented generations (RAG) to audit workflows. *Journal of Emerging Technologies in Accounting, 23*(1), 189–198. https://doi.org/10.2308/JETA-2024-041
25. Zhao, P., Zhang, H., Yu, Q., Wang, Z., Geng, Y., Yang, L., Zhang, W., Jiang, J., & Cui, B. (2024). Retrieval-augmented generation for AI-generated content: A survey. *arXiv*. https://doi.org/10.48550/arXiv.2402.19473