Journal of
Information Technology,
Cybersecurity, and Artificial Intelligence

(RESEARCH)

A Thematic Qualitative Literature Review of Risk Controls, Compliance Frameworks, and Auditability in Enterprise GRC

Dr. Ian Herzing
Associate Professor
Herzing University

Journal of Information Technology, Cybersecurity, and Artificial Intelligence, 2026, 3(3), 63-74

Article DOI: https://doi.org/10.70715/jitcai.2026.v3.i3.075

Abstract

This structured qualitative literature review examines how generative artificial intelligence (GenAI) governance is operationalized in enterprise governance, risk, and compliance (GRC). The analysis covered 37 studies, standards, regulatory materials, scholarly preprints, and professional guidance documents, with substantive GenAI governance and control literature concentrated in 2019-2026. Using an adapted PICO search framework and reflexive thematic analysis, this synthesis identified four themes: framework convergence without control specificity, lifecycle control as the dominant operational model, auditability as an evidence problem, and persistent gaps in accountability and continuous monitoring. Results suggest that modern GenAI GRC has moved beyond principles but remains immature as an assurance discipline because organizations still lack standardized evidence artifacts, control ownership, and continuous testing practices. This study contributes a control-oriented synthesis linking governance frameworks to auditable evidence artifacts and clarifies how enterprise teams can translate legal, standards, security, and audit guidance into practical control work.

Keywords: generative AI; GRC; AI auditability; AI risk management; compliance frameworks; internal control

1. Introduction

Generative artificial intelligence (GenAI), particularly large language models (LLMs), emerged as a major governance, risk, and compliance (GRC) concern because it combined rapid enterprise adoption with risks that were difficult to evaluate using traditional information technology control models. LLMs were adaptable across many downstream tasks, which made their risk profile highly dependent on context, deployment environment, user behavior, and organizational control maturity [1]. Prior research identified risks including discrimination, privacy leakage, misinformation, malicious use, unsafe human-computer interaction, and broader automation harms [2]. These risks were not solely technical; they also raised questions of accountability, assurance, documentation, legal compliance, and organizational responsibility.

AI governance therefore became an increasingly important field of inquiry. At the organizational level, AI governance was defined as the rules, practices, processes, and technical mechanisms used to align AI use with organizational strategy, values, legal requirements, and ethical principles [3]. Recent reviews showed that many AI governance approaches remained fragmented, with uncertainty about who was accountable, what was governed, when governance occurred in the AI lifecycle, and how governance was implemented through frameworks, tools, policies, and controls [4], [5]. This gap was especially visible in enterprise GRC, where organizations needed to translate high-level principles into auditable control activities, evidence artifacts, monitoring routines, and management-level oversight.

The regulatory and standards environment also accelerated during this period. NIST released the AI Risk Management Framework in 2023 and a Generative AI Profile in 2024 to help organizations identify and manage GenAI-specific risks [6], [7]. ISO/IEC 42001:2023 introduced a management-system standard for establishing and improving an AI management system [8]. The EU AI Act entered into force in 2024 and established a risk-based legal framework for AI systems and general-purpose AI models [9], [10]. Practitioner frameworks, including the OWASP Top 10 for LLM Applications and the Cloud Security Alliance AI Controls Matrix, translated GenAI risks into control categories and control objectives [11], [12].

Copyright © 2026 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution Liscense 4.0.

DOI: https://doi.org/10.70715/jitcai.2026.v3.i3.075

Despite this growth in guidance, operationalization remains underdeveloped. AI auditing scholarship shows that assurance practices contribute guidance, methods, tools, and stakeholder awareness, but that standards for auditability and evidence remain unsettled [13], [14]. Mökander et al. [1] argue that LLM auditing requires coordinated governance, model, and application audits rather than isolated technical testing. Accordingly, this review examines how GenAI governance is operationalized in enterprise GRC. The guiding research question is: How do organizations translate emerging GenAI governance frameworks into practical GRC controls, and where do gaps remain in auditability, accountability, and continuous compliance? The contribution is a GRC-facing synthesis that organizes fragmented legal, standards, cybersecurity, audit, and AI-governance literature into practical control-and-evidence vocabulary for enterprise risk, compliance, and assurance work. Figure 1 presents a layered operating model for enterprise GenAI governance, risk, and compliance that organizes the reviewed literature into three audit layers, governance, model, and application, supported by a lifecycle control band (design, develop, deploy, monitor, retire) and a set of auditable evidence artifacts, and it serves as the conceptual scaffold for the themes developed in Section 3 and the operating-model recommendations in Section 4.

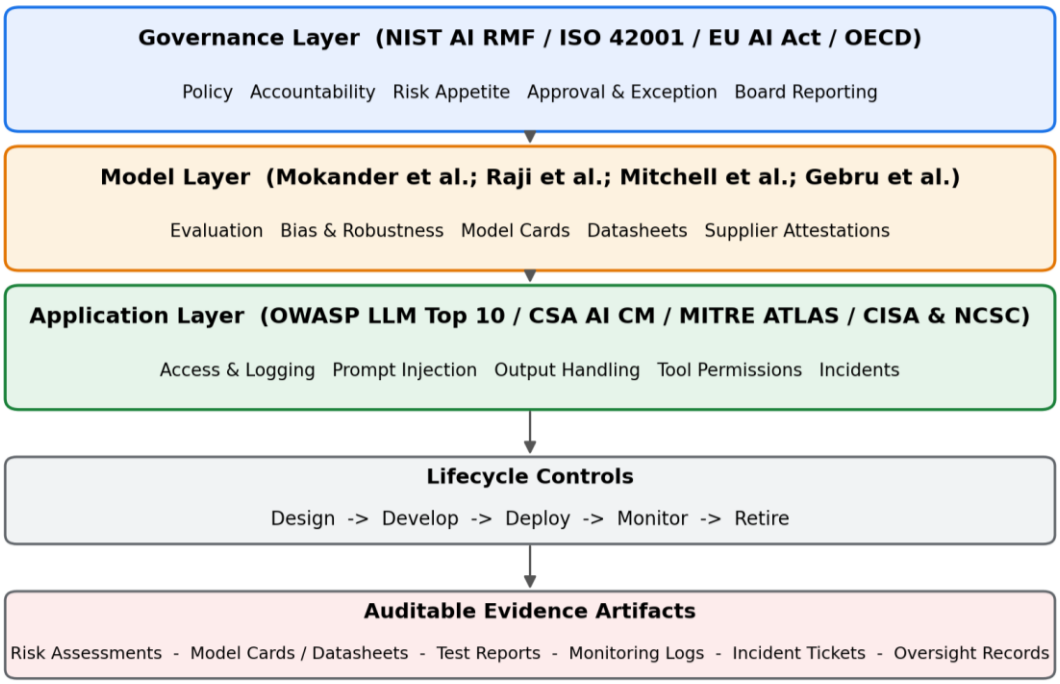

Figure 1: A layered GenAI GRC operating model

2. Methodology

2.1. Design

This study used a structured qualitative thematic literature review design supported by an adapted PICO framework and reflexive thematic analysis [15]. The design was appropriate because the purpose of the review was not to estimate effect size or produce an exhaustive systematic review, but to synthesize how academic and practitioner sources conceptualize and operationalize GenAI governance, control design, compliance alignment, and audit evidence. The analysis therefore emphasizes transparent search concepts, explicit eligibility criteria, source-type coding, and thematic synthesis across a curated corpus.

2.2. PICO Framework

A PICO-style framework was used to structure the analytic question and search strategy. Because the study examined organizational governance rather than a clinical intervention, the comparison element was treated as a contextual comparison across source types, frameworks, and traditional versus AI-specific control models. Table 1 summarized the adapted PICO framework.

Table 1: Adapted PICO Framework for the Review

PICO element	Operational definition	Search concepts	Inclusion focus

DOI: https://doi.org/10.70715/jitcai.2026.v3.i3.075

Population / Problem	Enterprises, public agencies, and regulated organizations adopting or governing GenAI, LLMs, foundation models, or AI-enabled applications.	enterprise OR organization* OR "public sector" OR regulated OR "large language model*" OR LLM OR "generative AI"	Organizational AI use, governance responsibility, or enterprise risk and compliance need.
Intervention / Phenomenon	Governance, risk management, compliance controls, auditability mechanisms, assurance practices, documentation, and AI management systems.	governance OR "risk management" OR compliance OR GRC OR audit* OR assurance OR control* OR documentation	Practical mechanisms for governing, controlling, assuring, or documenting AI systems.
Comparison / Context	AI-specific frameworks compared with traditional IT/cybersecurity GRC frameworks; peer-reviewed research compared with grey literature.	"NIST AI RMF" OR "ISO 42001" OR "EU AI Act" OR COBIT OR "NIST CSF" OR "OWASP LLM"	Evidence was coded for AI-specific governance, conventional GRC adaptation, or cross-framework mapping.
Outcome	Operationalized controls, accountable ownership, audit evidence, lifecycle monitoring, compliance readiness, control maturity, and unresolved gaps.	auditability OR "audit evidence" OR accountability OR monitoring OR "third-party risk" OR "model risk" OR "control maturity"	Contribution to understanding how GenAI governance became auditable, measurable, or operational.

Note. PICO was adapted for a qualitative GRC literature review; comparison referred to framework, source, and context comparisons rather than treatment-control comparisons.

2.3. Search Strategy

The search strategy was structured for replication across three primary academic databases: Scopus, Web of Science Core Collection, and IEEE Xplore. Supplementary discovery used ACM Digital Library, ScienceDirect, SpringerLink, arXiv, SSRN, and official grey-literature repositories because GenAI governance research develops rapidly and relevant work often appears first as preprints, standards, regulator materials, or professional guidance. Search terms combined GenAI or LLM concepts with governance, risk, compliance, auditability, assurance, documentation, and control concepts. The corpus spans foundational method and AI documentation sources from 2006 onward, with substantive GenAI governance, auditability, and GRC sources concentrated in 2019-2026. The strategy was intended to identify a defensible and diverse corpus for thematic synthesis, not to claim exhaustive database coverage.

Table 2: Search Strategy by Database and Source Type

Source / database	Role	Example search string	Rationale / output
Scopus	Primary academic database	("generative AI" OR "large language model*" OR LLM) AND (governance OR "risk management" OR compliance OR audit* OR control*)	Interdisciplinary research in information systems, cybersecurity, governance, and management.
Web of Science Core Collection	Primary academic database	("AI governance" OR "responsible AI governance" OR "AI auditing") AND (auditability OR compliance OR controls)	Indexed journal literature and systematic reviews on AI governance and auditability.
IEEE Xplore	Primary technical database	("large language model*" OR "generative AI") AND (security OR assurance OR governance OR audit* OR compliance)	Technical governance, assurance, cybersecurity, and control-oriented work.
ACM, ScienceDirect, SpringerLink	Supplementary databases	("AI governance" OR "LLM governance" OR "AI auditing") AND ("enterprise" OR "organization*" OR "risk")	FAccT, information systems, and AI ethics literature relevant to control design.
arXiv and SSRN	Preprint databases	("generative AI" OR LLM) AND (governance OR auditability OR compliance OR GRC)	Emerging auditability and control scholarship published before journal placement.

NIST, ISO, EU, OECD	Standards and regulatory grey literature	"AI RMF" OR "Generative AI Profile" OR "ISO/IEC 42001" OR "EU AI Act" OR "OECD AI Principles"	Authoritative governance, standards, and compliance frameworks.
OWASP, CSA, CISA/NCSC, ENISA, MITRE, COSO, ISACA, IIA	Professional grey literature	("LLM Top 10" OR "AI Controls Matrix" OR "secure AI" OR "AI audit toolkit" OR "internal control" OR ATLAS)	Control catalogs, audit tools, secure lifecycle guidance, and threat taxonomies.

2.4. Eligibility, Screening, and Analysis

Sources were eligible for inclusion if they addressed GenAI or LLM governance, AI risk management, AI compliance frameworks, AI auditability or assurance, enterprise control design, AI documentation artifacts, third-party or model risk management, or lifecycle monitoring of AI systems. Sources were excluded if they focused only on technical model performance without governance implications, discussed AI ethics at a purely philosophical level without operational controls, or addressed non-generative AI systems without transferable governance relevance.

The final corpus included 37 curated and citation-verified sources and intentionally mixed peer-reviewed research, scholarly preprints, regulatory documents, standards, and practitioner control frameworks. Each source was coded for source type, governance level, risk-control focus, compliance framework, audit evidence artifact, and implementation gap. Peer-reviewed studies were assessed for relevance, methodological transparency, and contribution to the research question, while grey literature was assessed for authority, currency, scope, and practical applicability. Citation verification checked that DOI, publisher, standards-body, regulator, arXiv, or official organizational records corresponded to the listed works and that each reference was used in the manuscript body or reference-corpus table.

Thematic analysis proceeded in six phases: familiarization with the corpus, initial coding, theme development, theme review, theme definition, and narrative synthesis [15]. Preliminary deductive codes included governance structure, accountability, AI inventory, risk classification, control ownership, documentation, audit evidence, human oversight, third-party/model risk, monitoring, incident response, regulatory mapping, and control maturity. Inductive codes were added as new concepts emerged during review.

Coding was performed manually by the author in Microsoft Excel rather than in dedicated qualitative analysis software (e.g., NVivo or ATLAS.ti), because the corpus size (n = 37) was tractable for hand coding and because mixing peer-reviewed research, standards, regulation, and practitioner guidance in one project makes proprietary qualitative software less useful than a flat, auditable codebook. A master coding matrix was built in Excel with one row per source and one column per coding dimension (source type, governance level, risk-control focus, compliance framework, audit evidence artifact, implementation gap, theme alignment, and notes). A second worksheet held the codebook itself: each code, its definition, its inclusion and exclusion rules, an exemplar passage, and the date the code was first introduced. A third worksheet held a coding memo log used to record interpretive decisions, code merges, code splits, and the trigger source for each inductive code. This three-worksheet structure provided a lightweight audit trail and allowed code application to be re-checked at any point in the synthesis.

The coding workflow combined deductive and inductive moves in a defined sequence. First, a deductive starter set was seeded from the dominant framework vocabulary in the corpus (NIST AI RMF Govern-Map-Measure-Manage, ISO/IEC 42001 management-system clauses, EU AI Act risk-tier obligations, OWASP LLM application risks, CSA AI Controls Matrix control families, and COSO internal control components). Each source was then read in full and tagged against this starter set in the matrix. Second, inductive codes were added whenever a passage carried governance meaning that the deductive set could not absorb without distortion - for example, evidence-by-design, control-owner ambiguity, supplier opacity, retrospective-audit infeasibility, and parallel-governance-committee risk all emerged inductively. Third, after the entire corpus had been coded, codes were grouped by conceptual proximity, competing groupings were tested against the source evidence, and the groupings that survived the most challenges became candidate themes. Fourth, candidate themes were defined, named, and stress-tested by checking whether every coded passage in the matrix still mapped to the theme without forcing. Where a theme did not fit, codes were re-allocated and the theme boundaries were rewritten. The final four themes reported in Section 3 are the product of this iterative deductive-then-inductive cycle, with the coded matrix retained as the underlying analytic artifact.

DOI: https://doi.org/10.70715/jitcai.2026.v3.i3.075

Table 3: Curated and Citation-Verified Reference Corpus and Use in the Review

Citation	Type	Primary contribution
[17] Arnold et al.	Peer-reviewed	Supplier declarations and assurance evidence.
[4] Batool et al.	Peer-reviewed	AI governance actors, artifacts, timing, and mechanisms.
[15] Braun and Clarke	Method	Thematic analysis method.
[18] Falco et al.	Peer-reviewed	Independent third-party audit model for AI assurance at enterprise scale.
[11] CSA	Grey	AI Controls Matrix and standards mapping.
[19] COSO	Grey	Internal control over GenAI.
[20] CISA and NCSC	Grey	Secure AI development lifecycle guidance.
[21] Dotan et al.	Preprint	NIST AI RMF maturity model.
[23] ENISA	Grey	AI cybersecurity practice framework.
[9] European Commission	Grey	EU AI Act implementation context.
[22] European Parliament and Council	Regulation	DORA operational resilience and third-party risk.
[10] European Parliament and Council	Regulation	EU AI Act legal requirements.
[24] Gebru et al.	Peer-reviewed	Dataset documentation as governance evidence.
[25] IIA	Grey	Internal audit framework for AI.
[8] ISO	Standard	AI management system requirements.
[26] ISACA	Grey	AI audit toolkit and control assessment.
[27] KPMG	Grey	Practitioner interpretation of COSO GenAI controls.
[13] Laine et al.	Peer-reviewed	Ethics-based AI auditing review.
[3] Mäntymäki et al.	Peer-reviewed	Organizational AI governance definition.
[28] McIntosh et al.	Peer-reviewed	GRC framework comparison for LLM commercialization.
[29] MITRE	Grey	AI threat matrix and adversarial technique taxonomy.
[30] Mitchell et al.	Peer-reviewed	Model cards as transparency artifacts.
[31] Mökander and Floridi	Peer-reviewed	Ethics-based auditing case study.
[1] Mökander et al.	Peer-reviewed	Three-layer LLM audit model.
[6] NIST	Grey	AI RMF Govern, Map, Measure, Manage model.
[7] NIST	Grey	GenAI-specific risk profile.
[32] NIST	Grey	Cybersecurity Framework 2.0 governance function.
[33] OECD	Grey	Trustworthy AI principles.
[12] OWASP Foundation	Grey	LLM application security risks.
[16] Page et al.	Method	PRISMA 2020 reporting guidance.
[5] Papagiannidis et al.	Peer-reviewed	Responsible AI governance review and framework.
[34] Raji et al.	Conference	Internal algorithmic auditing lifecycle.
[14] Schiff et al.	Peer-reviewed	Empirical AI ethics auditing ecosystem study.
[35] Schuett	Peer-reviewed	Risk management in the EU AI Act.
[36] SEC	Regulator	Cyber governance and disclosure model.

[37] Waltersdorfer et al.	Preprint	Continuous AI auditing infrastructure.
[2] Weidinger et al.	Preprint	Language model risk taxonomy.

3. Results

The first analytic cycle identified four themes. Table 4 summarized the themes, and the following subsections described how the evidence converged across academic research, standards, regulation, and professional guidance. Across themes, the evidence was strongest for framework convergence and control vocabulary, and weaker for empirical claims about which control combinations reduced risk in deployed enterprise settings.

Table 4: Summary of Themes

Theme	Label	Finding	Representative sources
1	Framework convergence without control specificity	Frameworks agreed on risk-based AI governance but left organizations to translate principles into testable controls.	[6], [7], [8], [10], [28]
2	Lifecycle control	Operational guidance placed controls across design, development, deployment, monitoring, third-party risk, and incident response.	[11], [12], [20], [23], [29]
3	Audit evidence	Auditability depended on documentation artifacts, control testing, supplier declarations, and assurance workpapers.	[17], [24], [30], [34], [25], [26]
4	Accountability and continuous compliance gaps	Organizations still lacked mature ownership models, monitoring infrastructure, and stable audit success measures.	[3], [14], [21], [37], [19]

3.1. Theme 1: Framework Convergence Without Control Specificity

The first theme was that the GenAI governance landscape showed broad agreement on risk-based governance, but less agreement on control specificity. NIST's AI RMF organized risk activity around Govern, Map, Measure, and Manage, while its GenAI profile added risks such as information integrity, misuse, and data confidentiality [6], [7]. ISO/IEC 42001 framed AI governance as a management system, and the EU AI Act required risk management, documentation, human oversight, and post-market monitoring for higher-risk systems [8], [10], [35]. OECD principles added normative expectations for trustworthy AI, while the European Commission [9] emphasized staged implementation of the AI Act [33]. However, the evidence suggested that these frameworks required translation before they could function as enterprise controls. McIntosh et al. [28] found that ISO 42001 was the most AI-specific framework among selected GRC frameworks, while COBIT-style governance was still useful for compliance alignment. Batool et al. [4] and Papagiannidis et al. [5] similarly found that AI governance research had advanced conceptually faster than operationally.

3.2. Theme 2: Lifecycle Control Became the Dominant Operational Model

The second theme was that operational guidance converged on lifecycle control. Rather than treating GenAI governance as a one-time model approval, the reviewed sources placed controls across design, development, deployment, monitoring, and retirement. Secure AI guidance from CISA and NCSC [20] and ENISA [23] framed security as a lifecycle issue that included secure design, supply chain risk, deployment hardening, and ongoing operation. OWASP Foundation [12] focused on LLM application risks such as prompt injection, sensitive information disclosure, supply chain vulnerabilities, excessive agency, and unbounded consumption. MITRE [29] contributed an adversarial threat taxonomy for AI systems, while CSA [11] translated governance and security expectations into 243 AI control objectives mapped to standards. The theme also extended beyond AI-specific documents. DORA treated ICT third-party resilience as a regulatory discipline, SEC disclosure rules required governance and risk-management transparency, and NIST CSF 2.0 made governance an explicit cybersecurity function [22], [32], [36]. These adjacent GRC models showed how GenAI controls could be linked to enterprise risk management, vendor oversight, and incident disclosure.

3.3. Theme 3: Auditability Depended on Documentation and Evidence Artifacts

The third theme was that auditability depended less on abstract principles than on durable evidence artifacts. Model cards, datasheets, and FactSheets showed how model behavior, data provenance, and supplier claims could be documented for review [17], [24], [30]. Raji et al. [34] described internal algorithmic auditing as an end-to-end process, while Mökander and Floridi [31] showed how ethics-based auditing could be used inside an enterprise setting. The LLM-specific literature extended this logic. Mökander et al. [1] argued that LLM audits required governance, model, and application layers, and Laine et al. [13] found that AI auditing produced knowledge contributions for multiple stakeholders. Professional audit sources also moved toward evidence. ISACA [26] described an AI audit toolkit for control assessment across the AI lifecycle, and the IIA [25] located internal audit within governance, management, and assurance domains. Falco et al. [18] argued that independent third-party audits, modeled on financial and safety audit

regimes, are necessary to give enterprise stakeholders durable assurance about AI system behavior, and they outlined the institutional preconditions for such audits to operate at scale.

3.4. Theme 4: Persistent Gaps Remained in Accountability and Continuous Compliance

The fourth theme was that organizations still lacked mature mechanisms for accountability and continuous compliance. Mäntymäki et al. [3] defined organizational AI governance broadly, but the corpus showed that definitions did not automatically produce ownership, testing cadence, or usable evidence. Schiff et al. [14] found that AI ethics auditing practice remained ambiguous, with variation in scope, stakeholder participation, and measurement of audit success. Dotan et al. [21] argued that organizations often lagged behind recommended AI risk-management practices, creating the risk of superficial governance. Waltersdorfer et al. [37] responded to the one-off audit problem by proposing infrastructure for continuous AI auditing. COSO [19] translated internal control concepts into GenAI control expectations, including monitoring, validation, and control mapping, while KPMG [27] provided a practitioner summary of that COSO guidance. The results suggested that continuous control monitoring was becoming central to GenAI GRC, but had not yet matured into a common assurance practice.

4. Discussion

4.1. Interpretation

This analysis found that modern GenAI GRC has moved from principles toward control-oriented practice, but that operational maturity remains uneven. The literature does not suggest that organizations lack frameworks; instead, it suggests that they lack stable translation mechanisms. NIST, ISO, OECD, and the EU provide governance expectations, while OWASP, CSA, MITRE, CISA/NCSC, ENISA, ISACA, IIA, and COSO provide more operational control language. The central GRC challenge is therefore integration: organizations must map normative obligations to policy owners, control objectives, evidence artifacts, testing procedures, exception handling, and board-level reporting. Because most reviewed sources are conceptual, standards-based, or practice-oriented rather than empirical evaluations of deployed controls, the findings should be read as a synthesis of current control logic rather than proof that any specific control bundle reduces GenAI risk.

4.2. Implications for Enterprise GRC

The findings suggest three practical implications. First, GenAI inventories should be treated as control objects rather than informal technology lists. Each GenAI use case needs an owner, purpose, data boundary, model or vendor source, risk classification, and evidence record. Second, organizations should separate governance controls from model and application controls. Governance controls address policy, accountability, approval, risk acceptance, and training; model controls address evaluation, documentation, bias, robustness, and provenance; application controls address access, logging, prompt injection, tool permissions, output handling, and incident response. Third, audit readiness depends on evidence-by-design. Evidence should include model cards, datasheets, supplier attestations, risk assessments, test reports, monitoring logs, human oversight records, vendor due diligence, and exception approvals. These implications also mean that GenAI governance should be embedded into ordinary GRC rhythms. Risk committees, change advisory boards, procurement reviews, cybersecurity assessments, privacy impact assessments, and internal audit planning should not treat GenAI as a special project outside the control environment. Instead, GenAI risks should be routed through existing approval, monitoring, issue-management, and reporting processes, with AI-specific evidence added where traditional controls are insufficient. For example, a vendor review can include standard security attestations, but also request model documentation, training-data constraints, incident notification commitments, and contractual rights to review material model changes. Similarly, a compliance review can map the same use case to AI Act obligations, privacy rules, cybersecurity requirements, and internal risk appetite statements. This integration matters because fragmented GenAI governance can create parallel committees, inconsistent approvals, and evidence gaps that weaken audit readiness. It also helps senior leaders see GenAI exposure as part of enterprise risk, rather than as a disconnected technology concern owned only by data science or security teams.

4.3. Toward an Enterprise GRC Operating Model

The synthesis also suggests that GenAI governance works best when it is organized as an operating model rather than as a standalone AI ethics statement. In this model, the enterprise first needs an authoritative GenAI inventory that classifies systems by business purpose, data exposure, model source, user group, autonomy level, legal context, and downstream decision impact. This inventory connects the governance expectations in NIST's AI RMF and ISO/IEC 42001 to concrete risk ownership because each use case can be linked to a control owner, evidence owner, business approver, and escalation pathway [6], [8]. Second, risk tiering should be used as a routing mechanism. Low-risk internal productivity uses may require acceptable-use controls, training, data-handling restrictions, and basic logging, while higher-risk uses involving customer decisions, regulated data, external outputs, or automated actions require formal

impact assessment, model or vendor review, human oversight, red-team testing, and post-deployment monitoring [7], [10]. Third, the operating model needs a policy-to-control mapping layer. This layer translates broad obligations, such as transparency, accountability, security, resilience, and human oversight, into testable control objectives and named evidence artifacts. That mapping is important because the reviewed frameworks do not use identical language. NIST emphasizes Govern, Map, Measure, and Manage; ISO emphasizes a management-system cycle; OWASP emphasizes application threats; CSA provides a control matrix; and COSO connects GenAI to internal control principles ([11] CSA; [19] COSO; [12] OWASP Foundation). A practical GRC program therefore needs crosswalks that make these frameworks mutually usable inside policy, risk, compliance, security, procurement, and audit workflows.

4.4. Audit and Compliance Practice Agenda

The findings further suggest that auditability must be designed before deployment rather than reconstructed afterward. Traditional audits often rely on stable systems, defined change windows, and durable transaction evidence, but GenAI systems introduce probabilistic outputs, prompt-level variability, vendor opacity, and continuing model updates. For that reason, audit workpapers need to capture both governance evidence and technical evidence. Governance evidence includes policy approvals, risk acceptance records, use-case inventories, training completion, third-party due diligence, exception logs, committee minutes, and board reporting. Technical evidence includes model cards, datasheets, test cases, prompt-injection testing, evaluation metrics, access logs, monitoring alerts, output review samples, incident tickets, and change records [17], [24], [30], [34]. This dual evidence requirement shows why internal audit, compliance, cybersecurity, privacy, procurement, and model-risk teams need a shared evidence taxonomy. Without one, each function can ask for different artifacts, creating fragmented assurance and duplicative reviews.

4.5. Comparative Synthesis Across Framework Families

A second-cycle reading of the four themes makes the comparative structure of the corpus easier to see. Sovereign and standards bodies (NIST, ISO, OECD, EU) anchor the governance layer with risk-based obligations but stop short of naming testable controls. Security-oriented bodies (OWASP, CSA, MITRE, CISA and NCSC, ENISA) populate the application layer with named threats, control objectives, and adversarial taxonomies but treat board-level accountability only in passing. Audit-oriented bodies (ISACA, IIA, COSO, and the algorithmic auditing scholarship of Raji et al. and Mokander et al.) sit between the two, translating obligations into evidence artifacts, ownership, and lifecycle checkpoints. Read this way, the corpus does not contain a single dominant framework but a triangle: governance frameworks supply the why, security frameworks supply the what, and audit frameworks supply the how. The operational gap exposed by the synthesis is therefore not a missing framework but a missing connector - the policy-to-control mapping layer described in Section 4.2 - that lets an organization read one obligation and answer all three questions about it in the same workflow. This connector argument is the load-bearing contribution of the present analysis and frames the future research agenda set out in Section 5.

4.6. Limitations

This study had several limitations that affected the strength, scope, and transferability of the synthesis. The most important limitation was that the screening, coding, theme development, and final interpretation were conducted by a single coder. Although a single-coder approach was practical for a working literature review and supported rapid synthesis of a fast-moving topic, it limited the dependability of the qualitative analysis because no second coder independently reviewed the corpus, challenged coding decisions, or tested the stability of the themes. Thematic analysis required interpretive judgment, especially when evidence crossed technical, legal, audit, and management domains [15]. As a result, the themes reflected a transparent but still researcher-mediated reading of the literature rather than an intercoder-validated consensus. A stronger design would have included at least one additional coder, double coding of a sample of studies, reconciliation meetings, a coding memo trail, and reporting of agreement or disagreement patterns before final theme naming.

A second limitation was that the review used a curated, citation-verified corpus rather than a completed exhaustive systematic search. The method followed an adapted PICO structure and drew from Scopus, Web of Science, IEEE Xplore, and official grey-literature repositories, but it did not report database-level record counts, deduplication totals, title-and-abstract exclusion counts, full-text exclusion reasons, or a PRISMA flow diagram. This limits reproducibility because another reviewer can follow the search concepts and source strategy but cannot fully reconstruct each screening decision from a complete exported record set [16]. The current synthesis should therefore be interpreted as a structured thematic study designed to support topic development, source mapping, and early conceptual synthesis, rather than as a final systematic review with exhaustive coverage claims.

A third limitation came from the rapid pace and uneven maturity of GenAI governance evidence. Several sources appeared as standards, regulator materials, professional guidance, preprints, audit toolkits, or practitioner control

catalogs rather than peer-reviewed empirical studies. Including grey literature was necessary because enterprise GRC practice was being shaped by NIST, ISO, the EU AI Act, OWASP, CSA, ISACA, COSO, and similar bodies faster than journal publication cycles could respond. However, this decision introduced unevenness in evidentiary weight. A peer-reviewed study, a legal instrument, a voluntary standard, and a practitioner checklist did not make the same kind of knowledge claim, and they could not be quality-appraised using one simple rubric. The synthesis partially addressed this by coding source type and contribution, but future work should apply separate appraisal criteria for empirical research, standards, regulation, technical control frameworks, and professional audit guidance.

A fourth limitation was conceptual scope. The analysis focused on operationalizing GenAI governance for enterprise GRC, with emphasis on risk controls, compliance frameworks, and auditability. This focus made the review useful for internal control, compliance, security, and assurance audiences, but it also meant that some adjacent literatures received less attention. For example, the review did not deeply synthesize model evaluation science, human factors research, labor impacts, environmental sustainability, procurement law, sector-specific model risk management, or public-sector administrative accountability. It also treated enterprise GRC as the primary implementation context, so findings may transfer less directly to small organizations, open-source communities, consumer AI products, or public agencies with different accountability structures.

Finally, this study was limited by the emerging nature of auditability itself. Many sources called for documentation, monitoring, evidence artifacts, human oversight, risk assessments, incident response, and supplier assurance, but fewer studies evaluated whether these controls actually reduced GenAI risk in deployed enterprise environments. The literature therefore supported a strong conceptual case for evidence-by-design, but it provided less empirical evidence about which controls worked best, how frequently controls should be tested, what evidence was sufficient for auditors, or how organizations should measure GenAI control maturity over time. Future research should test the proposed governance, model, and application control layers through case studies, audit simulations, and longitudinal studies of organizations implementing AI management systems.

5. Conclusion and Future Work

The reviewed literature shows that operationalizing GenAI governance requires more than selecting an AI framework. Enterprise GRC teams need a control architecture that connects legal requirements, management-system expectations, security threats, audit evidence, and continuous monitoring. The most promising direction is a layered model in which governance, model, and application audits are supported by lifecycle controls and evidence artifacts. The main unresolved gap is not the absence of guidance but the lack of consistent, auditable, repeatable, and empirically tested implementation practices across organizations. The synthesis therefore advances three contributions: a layered audit model that connects governance, model, and application controls; a comparative reading that locates each framework family in that model; and an evidence-by-design stance that treats auditability as a deployment-time design decision rather than a retrospective reconstruction. Together these contributions reframe enterprise GenAI GRC from a question of which framework to adopt to a question of how the available frameworks are stitched together into a single, testable control architecture.

The analysis points to several priorities for future research and practice. First, the field needs empirical studies of GenAI controls in real organizations. Many sources recommend governance, monitoring, documentation, and assurance, but few studies test whether particular control combinations reduce operational, legal, privacy, cybersecurity, or reputational risk. Case studies could compare organizations that implement ISO/IEC 42001-style management systems, NIST AI RMF mappings, OWASP LLM application controls, or CSA control mappings to identify what improves evidence quality and control maturity [6], [8], [11], [12]. Second, the field needs better measures of audit success. Schiff et al. [14] found ambiguity in AI ethics auditing, and Waltersdorfer et al. [37] showed the need for infrastructure supporting continuous audit. In enterprise GRC, audit success cannot be limited to whether a checklist exists; it must also include whether controls are owned, tested, remediated, monitored, and reported. Third, future work needs more attention to supplier and third-party assurance. Enterprises often use external foundation models, embedded AI features, and software vendors, so assurance depends on supplier transparency and contractual evidence as much as internal policy [17], [18]. Finally, practice needs clearer handoffs between governance, model, and application audits. The LLM audit literature shows that these layers are related but distinct; mature GRC programs therefore need to specify what each assurance layer tests, what evidence it produces, and how findings move into remediation and management reporting [1], [13].

6. References

[1] J. Mökander, J. Schuett, H. R. Kirk, and L. Floridi, "Auditing large language models: A three-layered approach," AI and Ethics, vol. 4, pp. 1085-1115, 2024, doi: 10.1007/s43681-023-00289-2.

[2] L. Weidinger et al., "Ethical and social risks of harm from language models," arXiv:2112.04359, 2021, doi: 10.48550/arXiv.2112.04359.

[3] M. Mäntymäki, M. Minkkinen, T. Birkstedt, and M. Viljanen, "Defining organizational AI governance," AI and Ethics, vol. 2, pp. 603-609, 2022, doi: 10.1007/s43681-022-00143-x.

[4] A. Batool, D. Zowghi, and M. Bano, "AI governance: A systematic literature review," AI and Ethics, vol. 5, pp. 3265-3279, 2025, doi: 10.1007/s43681-024-00653-w.

[5] E. Papagiannidis, P. Mikalef, and K. Conboy, "Responsible artificial intelligence governance: A review and research framework," The Journal of Strategic Information Systems, vol. 34, no. 2, Art. no. 101885, 2025, doi: 10.1016/j.jsis.2024.101885.

[6] National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," 2023, doi: 10.6028/NIST.AI.100-1.

[7] National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile," 2024, doi: 10.6028/NIST.AI.600-1.

[8] International Organization for Standardization, "ISO/IEC 42001:2023: Information technology - Artificial intelligence - Management system," 2023. [Online]. Available: https://www.iso.org/standard/81230.html

[9] European Commission, "AI Act enters into force," 2024. [Online]. Available: https://commission.europa.eu/news/ai-act-enters-force-2024-08-01_en

[10] European Parliament and Council, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence," Official Journal of the European Union, 2024. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[11] Cloud Security Alliance, "AI Controls Matrix," 2025. [Online]. Available: https://cloudsecurityalliance.org/artifacts/ai-controls-matrix

[12] OWASP Foundation, "OWASP Top 10 for Large Language Model Applications," 2025. [Online]. Available: https://owasp.org/www-project-top-10-for-large-language-model-applications/

[13] J. Laine, M. Minkkinen, and M. Mäntymäki, "Ethics-based AI auditing: A systematic literature review on conceptualizations of ethical principles and knowledge contributions to stakeholders," Information & Management, vol. 61, no. 5, Art. no. 103969, 2024, doi: 10.1016/j.im.2024.103969.

[14] D. S. Schiff, S. Kelley, and J. Camacho Ibáñez, "The emergence of artificial intelligence ethics auditing," Big Data & Society, vol. 11, no. 4, pp. 1-16, 2024, doi: 10.1177/20539517241299732.

[15] V. Braun and V. Clarke, "Using thematic analysis in psychology," Qualitative Research in Psychology, vol. 3, no. 2, pp. 77-101, 2006, doi: 10.1191/1478088706qp063oa.

[16] M. J. Page et al., "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," BMJ, vol. 372, Art. no. n71, 2021, doi: 10.1136/bmj.n71.

[17] M. Arnold et al., "FactSheets: Increasing trust in AI services through supplier declarations of conformity," IBM Journal of Research and Development, vol. 63, no. 4/5, pp. 6:1-6:13, 2019, doi: 10.1147/JRD.2019.2942288.

[18] G. Falco, B. Shneiderman, J. Badger, R. Carrier, A. Dahbura, D. Danks, M. Eling, A. Goodloe, J. Gupta, C. Hart, M. Jirotka, H. Johnson, C. LaPointe, A. J. Llorens, A. K. Mackworth, C. Maple, S. E. Palsson, F. Pasquale, A. Winfield, and Z. S. Yeong, "Governing AI safety through independent audits," Nature Machine Intelligence, vol. 3, no. 7, pp. 566-571, 2021, doi: 10.1038/s42256-021-00370-7.

[19] Committee of Sponsoring Organizations of the Treadway Commission, "Achieving effective internal control over generative AI," 2026. [Online]. Available: https://www.coso.org/generative-ai

[20] Cybersecurity and Infrastructure Security Agency and National Cyber Security Centre, "Guidelines for secure AI system development," 2023. [Online]. Available: https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development

[21] R. Dotan, B. Blili-Hamelin, R. Madhavan, J. Matthews, and J. Scarpino, "Evolving AI risk management: A maturity model based on the NIST AI Risk Management Framework," arXiv:2401.15229, 2024, doi: 10.48550/arXiv.2401.15229.

[22] European Parliament and Council, "Regulation (EU) 2022/2554 on digital operational resilience for the financial sector," Official Journal of the European Union, 2022. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32022R2554

[23] European Union Agency for Cybersecurity, "Multilayer framework for good cybersecurity practices for AI," 2023. [Online]. Available: https://www.enisa.europa.eu/publications/multilayer-framework-for-good-cybersecurity-practices-for-ai

[24] T. Gebru et al., "Datasheets for datasets," Communications of the ACM, vol. 64, no. 12, pp. 86-92, 2021, doi: 10.1145/3458723.

[25] Institute of Internal Auditors, "The IIA's Artificial Intelligence Auditing Framework," 2nd ed., 2024. [Online]. Available: https://www.theiia.org/en/content/tools/professional/2023/the-iias-updated-ai-auditing-framework

[26] ISACA, "Artificial Intelligence Audit Toolkit," 2024. [Online]. Available: https://www.isaca.org/resources/artificial-intelligence

[27] KPMG, "COSO releases roadmap on internal control over generative AI," 2026. [Online]. Available: https://kpmg.com/us/en/frv/reference-library/2026/coso-releases-roadmap-internal-control-over-generative-ai.html

[28] T. R. McIntosh et al., "From COBIT to ISO 42001: Evaluating cybersecurity frameworks for opportunities, risks, and regulatory compliance in commercializing large language models," Computers & Security, vol. 144, Art. no. 103964, 2024, doi: 10.1016/j.cose.2024.103964.

[29] MITRE, "MITRE ATLAS," n.d. [Online]. Available: https://atlas.mitre.org/

[30] M. Mitchell et al., "Model cards for model reporting," in Proc. Conf. Fairness, Accountability, and Transparency, 2019, pp. 220-229, doi: 10.1145/3287560.3287596.

[31] J. Mökander and L. Floridi, "Operationalising AI governance through ethics-based auditing: An industry case study," AI and Ethics, vol. 3, pp. 451-468, 2023, doi: 10.1007/s43681-022-00171-7.

[32] National Institute of Standards and Technology, "The NIST Cybersecurity Framework (CSF) 2.0," 2024, doi: 10.6028/NIST.CSWP.29.

[33] OECD, "OECD AI Principles," 2024. [Online]. Available: https://www.oecd.org/en/topics/ai-principles.html

[34] I. D. Raji et al., "Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing," in Proc. 2020 Conf. Fairness, Accountability, and Transparency, 2020, pp. 33-44, doi: 10.1145/3351095.3372873.

[35] J. Schuett, "Risk management in the Artificial Intelligence Act," European Journal of Risk Regulation, vol. 15, no. 2, pp. 367-385, 2024, doi: 10.1017/err.2023.1.

[36] Securities and Exchange Commission, "Cybersecurity risk management, strategy, governance, and incident disclosure," 2023. [Online]. Available: https://www.sec.gov/rules-regulations/2023/07/s7-09-22

[37] L. Waltersdorfer, F. J. Ekaputra, T. Miksa, and M. Sabou, "AuditMAI: Towards an infrastructure for continuous AI auditing," arXiv:2406.14243, 2024, doi: 10.48550/arXiv.2406.14243.