RAIS RESEARCH ASSOCIATION for INTERDISCIPLINARY STUDIES
March 12-13, 2026
DOI:10.5281/zenodo.19553772

Evaluating Control-Based AI Governance in Cybersecurity GRC Programs: An Expert Assessment Study

Miranda Stanfield

Capitol Technology University, Laurel, MD, USA, miranda.stanfield@gmail.com

Abstract: This study proposes and evaluates the Lifecycle-Integrated AI Governance Control Framework (LIAGCF), a Design Science Research artifact that maps 16 administrative, technical, and operational AI governance controls across seven lifecycle phases. Each control aligns to NIST AI RMF core functions and NIST SP 800-53 control families within a unified enterprise cybersecurity GRC architecture. Structured expert review (N = 10) validated NIST AI RMF structural alignment (9 of 10) and confirmed SP 800-53 control family assignments as technically defensible (7 of 10). Lifecycle coverage was validated as comprehensive, and the framework was affirmed as a meaningful contribution to AI governance scholarship. Theoretically, the LIAGCF addresses a documented structural gap as the first lifecycle-integrated, control-categorized governance framework to operationalize both NIST AI RMF core functions and NIST SP 800-53 control families within a unified architecture, bridging the persistent disconnect between strategic AI risk oversight and operational cybersecurity enforcement. In practice, the LIAGCF equips GRC practitioners, CISOs, and federal risk management program leads with a standards-aligned, auditable governance structure that operationalizes AI risk management intent as assignable, lifecycle-anchored controls, with designated implementation artifact requirements.

Keywords: AI Governance, Cybersecurity Governance, Governance, Risk, and Compliance (GRC), Control-Based Governance, NIST AI Risk Management Framework, NIST SP 800-53, Auditability and Assurance, Expert Assessment, Lifecycle-Integrated AI Governance Control Framework (LIAGCF)

1. Introduction

1.1 Background

As artificial intelligence becomes embedded in high-stakes enterprise and federal decision-making environments, organizations are confronting a structural governance problem that existing institutional arrangements are poorly equipped to address. Empirical research has documented that AI adoption across corporate, public-sector, and critical infrastructure contexts has consistently outpaced the development of commensurate oversight mechanisms. Scholars identify the resulting governance gaps as rooted in algorithmic opacity, unmonitored performance degradation, and fragmented regulatory compliance (Egwuatu, 2025; Robles & Mallinson, 2026; Saurabh, 2026). Traditional information technology governance frameworks have proven insufficient to address these gaps because they were designed for deterministic, static systems rather than for probabilistic models that adapt, degrade, and distribute accountability across organizational boundaries in ways that defy ex-ante verification (Owolabi, 2026). While some have proposed placing AI governance in legal departments, human resources functions, or standalone ethics offices, these arrangements lack the institutional infrastructure needed to operationalize oversight. None has the control frameworks, risk registers, continuous monitoring pipelines, audit authority, or regulatory compliance architecture that governance at scale requires. By contrast, Cybersecurity Governance, Risk, and Compliance (GRC) programs already possess precisely these tools. GRC functions are institutionally positioned to conduct AI risk assessments, implement and test controls, manage findings through remediation workflows, and produce the evidentiary documentation that regulators increasingly demand (Robles & Mallinson, 2026; Saurabh, 2026). Despite this structural alignment, the academic literature has not yet empirically examined whether cybersecurity GRC practitioners recognize AI governance as within their

organizational mandate, nor whether existing GRC competencies are perceived as sufficient to meet the demands of AI deployment. This study addresses that gap.

1.2 Problem Statement

The accelerating integration of artificial intelligence into federal and enterprise information systems has outpaced the development of commensurate organizational governance infrastructure, producing significant and documented consequences. Empirical assessment of federal AI governance reveals that fewer than 40% of 45 legally mandated governance requirements have been verified as implemented across federal agencies, despite more than 1,700 documented AI use cases governmentwide (Lawrence et al., 2023). A significant portion of these are classified as rights- or safety-impacting, and deployment scale continues to exceed the governance frameworks necessary to ensure transparency, auditability, and accountability (Lawrence et al., 2023; Nwinyi & Oman-Amoako, 2026). The general problem is that AI systems are deployed at scale in environments where lifecycle-integrated governance infrastructure is insufficient to match the scope or complexity of deployment. Within cybersecurity programs, the problem is that GRC functions operating under NIST SP 800-53 and the NIST AI RMF lack an integrated, operationalizable governance model. No existing model addresses AI-specific risks across the full system lifecycle (Kabir et al., 2025; Rahman et al., 2026). While supplemental frameworks have been proposed in the literature, an empirically grounded expert evaluation of whether control-based governance aligned to these standards can adequately govern AI systems within enterprise cybersecurity GRC programs has not been conducted. This study addresses that gap through a structured expert assessment.

1.3 Research Purpose and Questions

The purpose of this study is to design and expert-validate the Lifecycle-Integrated AI Governance Control Framework (LIAGCF) as a control-based governance artifact for enterprise cybersecurity GRC programs. Assessment targets five dimensions: structural alignment, operationalizability, lifecycle coverage, theoretical contribution, and feasibility of operational integration through structured expert review.

Five research questions correspond to the primary evaluation areas of the Design Science Research artifact assessment.

RQ1: How effectively do the LIAGCF governance controls align with the four NIST AI RMF core functions, GOVERN, MAP, MEASURE, and MANAGE?

RQ2: How appropriately do the NIST SP 800-53 control family assignments operationalize each governance objective within the LIAGCF?

RQ3: How comprehensively does the LIAGCF represent governance requirements across the seven AI system lifecycle phases from Strategy and Design through Retirement, and what governance domains or risk considerations remain insufficiently addressed?

RQ4: Does the integration of NIST AI RMF functions and NIST SP 800-53 control families within a unified, lifecycle-aligned governance control structure represent a meaningful and distinctive contribution to AI governance practice and scholarship?

RQ5: What operational barriers would organizations encounter when integrating the LIAGCF into existing GRC infrastructure?

1.4 Scope and Delimitations

The LIAGCF is scoped to operationalizable governance controls within cybersecurity GRC programs. AI ethics constructs such as fairness and transparency, where these operate independently of verifiable control structures, fall outside its architectural scope. The adequacy or suitability of the NIST AI RMF and NIST SP 800-53 as governance standards is not assessed; both are treated as authoritative and assumed rather than examined. SP 800-53 control

configuration, testing, and system-level documentation fall under system-level security engineering; this study operates at the governance architecture level.

Operational effectiveness is not evaluated, as no deployment has been conducted. Expert validation supports face and construct validity within this study's evaluation design; evidence of operational effectiveness would require longitudinal deployment studies, identified as a direction for future research. The 16 governance controls represent a purposive cross-section selected to validate the framework's structural alignment logic, not a complete enterprise control catalog. Organizations adopting the LIAGCF are expected to extend the control set to reflect their specific risk profiles and operational contexts.

1.5 Significance of the Study

This study makes a dual contribution. Theoretically, the LIAGCF addresses a documented structural gap as the first lifecycle-integrated, control-categorized governance framework to operationalize both NIST AI RMF core functions and NIST SP 800-53 control families within a unified architecture, bridging the persistent disconnect between strategic AI risk oversight and operational cybersecurity enforcement. Practically, the LIAGCF equips GRC practitioners, CISOs, and federal risk management program leads with a standards-aligned, auditable governance structure that operationalizes AI risk management intent as assignable, lifecycle-anchored controls with designated implementation artifact requirements.

1.6 Organization of the Paper

Section 2 reviews the literature across five literature domains, culminating in a gap synthesis that identifies the specific structural deficits addressed by this study. Section 3 presents the theoretical framework underlying the LIAGCF, including its conceptual foundations, alignment pathway, and 16-control design logic. Section 4 describes the DSR methodology, expert-review strategy, and validity considerations. Section 5 reports findings organized by research question. Section 6 discusses interpretation, framework revisions, and study limitations. Section 7 presents conclusions, practitioner recommendations, and directions for future research.

2. Literature Review

2.1 AI Governance: Definitions, Frameworks, and the Control Gap

Dominant AI governance frameworks are structured at the level of principles and risk categories rather than on assignable, verifiable controls. The NIST AI RMF, ISO/IEC 42001, the EU AI Act, and the OECD AI Principles represent the dominant reference structures guiding organizational governance of AI systems (Finch & Butt, 2025; Paul Greene, 2025). Despite their breadth, each shares this common structural limitation.

Empirical and review-based scholarship corroborates this structural limitation. Gupta (2025) explicitly identifies the translation gap, documenting that the NIST AI RMF and ISO 42001 are positioned at high conceptual levels, providing principles without actionable controls. Gupta further documents that organizations lack systematic governance and risk management practices throughout the AI development lifecycle, as well as mechanisms to measure compliance quantitatively. Finch and Butt (2025) corroborate this through a systematic review documenting structural asymmetry between regulatory ambition and institutional implementation capacity, including irresolvable conflicts between the EU AI Act and GDPR that fragment enterprise accountability.

The structural limitations documented above carry measurable operational consequences. Essien et al. (2025) document a 26-percentage-point gap between enterprise AI deployment rates (68%) and GRC framework implementation maturity (42%) as of 2025. Casimir (2026) documents downstream consequences, including model hallucination, bias amplification, and escalating technical debt, outcomes traceable to the absence of structured

control assignments. Reuel and Undheim (2024) argue that AI governance must co-evolve with AI capability, as periodic disclosure regimes produce structural misalignment when applied to continuously evolving systems.

Existing approaches address components of this gap without producing an integrated control-level governance structure. Paul Greene (2025) proposes a "govern once/comply many" model that maps NIST AI RMF functions to the cybersecurity confidentiality, integrity, and availability (CIA) principles. Adebayo (2023) demonstrates measurable improvements in supply chain risk governance through NIST AI RMF and GRC alignment. Butt et al. (2025) embed policy enforcement into AI development and MLOps pipelines. Lestari et al. (2025) demonstrate verifiable improvements in governance maturity through the integration of multi-standard NIST AI RMF. None of these approaches produces a lifecycle-integrated, control-level governance structure mapping AI RMF functions to SP 800-53 control family assignments within an enterprise GRC architecture. This gap represents the specific problem this study addresses by designing and evaluating the LIAGCF.

2.2 Cybersecurity GRC Programs and NIST SP 800-53

Although NIST Special Publication 800-53 was not developed to address AI-specific risks, its structured security control catalog has gained increasing recognition for its applicability to AI governance contexts. The fifth edition, published in 2020, expanded coverage to include privacy controls and introduced the Supply Chain Risk Management (SR) control family, both directly relevant to AI system governance (National Institute of Standards and Technology [NIST], 2020).

Empirical evidence supports this applicability across auditing, GRC platform integration, and AI-driven assurance contexts. Kemp (2025) documents that AI system auditing frameworks integrate SP 800-53 controls alongside the NIST AI RMF to assess adversarial robustness, privacy control compliance, and continuous monitoring requirements. Faruq's (2025) meta-analysis confirms that integrating the NIST framework into GRC platforms yields measurable improvements in audit readiness and compliance posture. Urhobo (2024) and Odedina (2024) document the structural transformation of GRC from periodic compliance functions to continuous assurance architectures. Jooda and Onukak (2023) demonstrate that AI-driven GRC systems achieve integrated risk modeling, compliance automation, and threat intelligence fusion when AI is coordinated with human oversight.

Among existing frameworks that integrate SP 800-53 with AI lifecycle governance, Rahman et al. (2026) produce the most directly relevant prior work. Their four-layer framework maps AI threat categories to SP 800-53 control families across seven AI lifecycle phases within a FedRAMP-aligned government cloud security architecture. The authors acknowledge, however, that the framework is scoped to the control family level and that operational implementation will require SP 800-53 control-level overlays consistent with SP 800-53B baseline guidance. The framework is bound to U.S. federal cloud infrastructure under FedRAMP authorization; enterprise GRC programs and commercial architectures fall outside its scope. No administrative, technical, and operational control categorization is applied as an organizing structure.

Neither Rahman et al. (2026) nor existing adjacent work produces an operationalized governance control structure organized by administrative, technical, and operational control category and integrated within an enterprise GRC program architecture. This structural absence represents the specific gap addressed by this study through the design and evaluation of the LIAGCF. Sudarsan et al. (2025) operationalize NIST AI RMF with evidence-driven controls across the AI lifecycle phases for critical infrastructure. They do not map to SP 800-53 control families or apply administrative, technical, and operational categorization, representing adjacent but structurally distinct work.

2.3 The AI System Lifecycle in Governance Contexts

Published AI lifecycle models have converged on a seven-phase governance structure that differs fundamentally from traditional software lifecycle frameworks. Analysis of models published between 2022 and 2026 confirms that Strategy and Design, Data Preparation, Development, Procurement, Deployment, Operations, and Retirement are the foundational phases for comprehensive AI system oversight. Accordingly, AI system governance has been reconceptualized as a lifecycle-integrated discipline rather than a static compliance function applied at discrete checkpoints.

The scholarly literature supports this seven-phase structure by documenting the governance properties that lifecycle-integrated frameworks must address. Leon (2026) establishes that lifecycle governance must account for AI systems' distinctive properties, including data dependencies, continuous learning, and probabilistic outputs, requiring governance structures capable of evolving alongside the AI system's operational state. Katakam (2025) exemplifies this through the Ethical Lifecycle Governance Framework (ELGF), embedding transparency, human oversight, validation, and continuous monitoring from data intake through retirement. Rashid et al. (2026) demonstrate that lifecycle-integrated NIST-aligned security architectures can simultaneously enforce governance, adversarial robustness, and regulatory compliance in critical infrastructure contexts. Saurabh (2026) documents that secure AI adoption models require lifecycle oversight, security-by-design principles, and continuous risk assessment across all phases.

Within this seven-phase model, the Retirement and Decommissioning phase remains the most systematically underspecified component of lifecycle governance. Frimpong (2026) documents that retired AI systems pose residual risks, including unrevoked model access, persistent data exposure, and residual inference capabilities, coining the term "AI debris" to describe this governance failure category. Obimakinde (2026) proposes a socio-technical decommissioning framework. Khalid and Jaffery (2025) document the governance requirements for system retirement and decommissioning, including secure data deletion, model parameter disposal, and audit documentation. Neither work integrates Retirement governance into a GRC control structure with SP 800-53 control family assignments. This study addresses that gap through the LIAGCF, which treats Retirement as a fully specified lifecycle phase with explicit administrative, technical, and operational control assignments within a NIST-aligned enterprise GRC architecture. That design decision distinguishes the LIAGCF from all adjacent prior work reviewed in this section.

2.4 Control-Based Governance Models for AI

The administrative, technical, and operational control taxonomy provides the theoretical and empirical foundation for the LIAGCF's structural design. Control-based governance models in information security distinguish three functional control categories: administrative controls, encompassing policies, governance structures, and organizational accountability; technical controls, implementing security mechanisms through automated or technological means; and operational controls, covering monitoring, incident response, and change management. Hong et al. (2003) establish the theoretical foundation through an integrated system theory identifying preventive, detective, and corrective control functions. Veiga and Eloff (2007) demonstrate its application as a governance discipline requiring formal, verifiable control mechanisms. Yaokumah (2017) provides empirical evidence that administrative access control policies significantly influence technical access control mechanisms through accountability structures, justifying this study's decision to maintain explicit categorical boundaries within the LIAGCF.

The application of this taxonomy to AI governance has been emergent and largely implicit. Catapang (2026) develops an ethics-by-design AI control architecture with a triple-gate structure, Metric gates, Governance gates, and Eco gates. This architecture functionally maps to technical, administrative, and operational dimensions without invoking the taxonomy

by name, representing the closest published AI governance design to the LIAGCF's structural approach in functional, if not terminological, terms. Hundeyin et al. (2025) provide the most explicit published application, mapping federated learning as a technical control, immutable logging and explainability as operational controls, and governance standards updates as administrative controls. Hermansyah and Zakaria (2024) integrate governance structures, technical controls, and agile practices with NIST SP 800-53 and NIST AI RMF, but do not systematically apply the taxonomy across the AI lifecycle phases. Across all three works, the taxonomy is present as a functional organizing principle but is never applied systematically as a named primary structure.

AI governance frameworks consistently converge on three-layer control architectures that are functionally equivalent to the administrative, technical, and operational taxonomy but do not apply it as a primary organizing structure. Saleh and Abdulsalam (2025) demonstrate that administrative controls translate into verifiable compliance artifacts within ISO/IEC 42001. This study addresses that gap through the LIAGCF, applying this taxonomy as the structural backbone for all 16 controls across seven lifecycle phases. Each control is aligned to a NIST AI RMF function and mapped to NIST SP 800-53 control families within a unified enterprise GRC architecture.

2.5 Design Science Research in Governance and Security Studies

Design Science Research provides the methodological foundation for the LIAGCF's artifact design and expert panel evaluation strategy. First formalized by Hevner et al. (2004), DSR provides a structured process for producing purposefully designed artifacts addressing defined organizational problems through iterative cycles of design, demonstration, and assessment. The six-activity process model established by Peffers et al. (2007) is widely adopted in cybersecurity governance research and supports expert panel evaluation without prescribing minimum sample sizes, which is directly relevant to the LIAGCF's expert panel structure.

Multiple recent studies confirm DSR's applicability to this research class. Melaku (2023) applies DSR to develop a cybersecurity governance framework evaluated through expert panel assessment, providing a direct methodological precedent. Yasin et al. (2020) apply the six-stage DSRM to produce COBIT 2019 and ISO 27001-aligned governance artifacts, demonstrating DSR's applicability to standards-aligned governance deliverables. Mwanje et al. (2023) apply DSR to produce an information security governance framework for SMEs evaluated through expert assessment. Wijaya et al. (2026) employ DSR grounded in COBIT 2019 to develop an Ambidextrous AI Governance Framework validated through expert interviews, document analysis, and governance maturity simulation. Ramli and Darus (2026) develop an AI-Ready ICT Security framework grounded in NIST CSF 2.0 with expert review and bounded pilot evaluation. These precedents collectively confirm that DSR with expert panel evaluation is the appropriate methodological choice for governance framework artifact development at the validation stage represented by this study.

The adequacy of the LIAGCF's expert panel size is grounded in Venable et al.'s (2016) Framework for Evaluation in Design Science Research (FEDS). FEDS establishes that small expert panels constitute appropriate naturalistic evaluation for governance artifacts at early validation stages, providing the methodological foundation for the LIAGCF's evaluation strategy. Taken together, these studies situate the LIAGCF's methodological approach within a well-documented tradition of governance artifact development and expert-based evaluation in cybersecurity governance research.

2.6 Summary of Literature Gaps

The design of the LIAGCF is grounded in five structurally distinct gaps, each documented by peer-reviewed evidence, that define the problem space this study addresses.

First, no peer-reviewed framework translates the NIST AI RMF's four core functions, GOVERN, MAP, MEASURE, and MANAGE, into assigned SP 800-53 control family mappings within a unified enterprise cybersecurity GRC program architecture. Rahman et al. (2026), the most directly adjacent work, maps at the control family level only and is scoped exclusively to U.S. federal cloud infrastructure under FedRAMP authorization. Second, no published AI governance model uses the categorization of administrative, technical, and operational control as its primary organizational structure, despite the taxonomy's established pedigree in information security governance. Third, no NIST-aligned GRC control structure integrates the Retirement and Decommissioning phase with SP 800-53 control family assignments, despite documented residual risks (Frimpong, 2026). Fourth, enterprise GRC programs lack a lifecycle-integrated AI governance control model for their operational context. Fifth, major principles-based frameworks from NIST, ISO, and the EU do not provide the operationalized, assignable control-level mappings required for auditable enterprise GRC program compliance (Gupta, 2025; Finch & Butt, 2025).

Taken together, these five gaps define a structural problem that no existing framework resolves. This study responds to all five through the LIAGCF, mapping 16 governance controls, four per NIST AI RMF core function, across seven AI lifecycle phases. SP 800-53 control family assignments are organized by administrative, technical, and operational categories within an enterprise GRC architecture. Section 3 presents the LIAGCF's theoretical foundations and design logic.

3. Theoretical Framework

3.1 Conceptual Foundations of the LIAGCF

The LIAGCF's structural design is grounded in three theoretical foundations: control theory, lifecycle governance theory, and dual-standards alignment. Applied to information governance, control theory holds that organizational intent must be operationalized through specific, verifiable mechanisms rather than articulated through principles alone (Hevner et al., 2004; March & Smith, 1995). In AI governance, this distinction is consequential. A policy declaring that AI systems will be used responsibly does not constitute governance unless it is translated into defined controls with identifiable owners, documented artifacts, and enforceable review processes. This study operationalizes this principle through the LIAGCF by assigning each governance objective a specific control category and a corresponding NIST SP 800-53 control family. Each objective also receives an example artifact that serves as implementation evidence, ensuring that governance intent is auditable (NIST, 2020).

The second foundation, lifecycle governance theory, holds that systems passing through distinct developmental and operational phases require phase-specific oversight mechanisms (NIST, 2023). A uniform control set applied across all phases will impose a disproportionate governance burden on low-risk phases while leaving critical transition points, particularly development-to-deployment and operations-to-retirement, exposed to governance gaps. The governance requirements for a model in data preparation differ structurally from those governing a deployed or retiring model (NIST, 2023). This study addresses this through the LIAGCF by assigning each control to an explicit AI lifecycle phase, enabling phase-gated governance reviews aligned with the governance and risk requirements of each lifecycle phase.

The third foundation addresses the structural disconnect between strategic risk oversight and operational enforcement (Hevner et al., 2004). Organizations often implement AI risk management programs without translating strategic alignment into the specific control structures required by their compliance programs. This study addresses this gap through dual-standards alignment in the LIAGCF. Each control is mapped simultaneously to a NIST AI RMF core function, positioning it within the strategic risk management cycle (NIST, 2023),

and to a NIST SP 800-53 control family, providing the operational enforcement pathway for implementing each governance objective (NIST, 2020).

3.2 The Six-Column Alignment Pathway

The six-column alignment pathway (Table A1) is the LIAGCF's primary translation mechanism, converting governance intent into verifiable implementation requirements across six sequential dimensions: Control Category, Control Objective, NIST AI RMF Function, SP 800-53 Control Family, AI Lifecycle Phase, and Governance Artifacts. Each column plays a distinct role in the governance translation chain. Together, they address what Gupta (2025) identifies as the persistent failure of existing frameworks: the inability to translate high-level principles into actionable, assignable controls.

The first column, Control Category, classifies each control as Administrative, Technical, or Operational, placing it within the taxonomy that information security governance literature has established as foundational to structured control architectures (Hong et al., 2003; Yaokumah, 2017). The second column, Control Objective, states the governance purpose each control is intended to achieve, ensuring this study's design intent is explicit and auditable. The third column assigns each control to one of the four NIST AI RMF core functions: GOVERN, MAP, MEASURE, or MANAGE. Paul Greene (2025) demonstrates that mapping controls to AI RMF functions enables a common compliance language, reducing the structural complexity that arises from managing multiple parallel governance obligations.

The fourth column assigns a specific SP 800-53 control family to each governance control, thereby providing the operationalization mechanism that existing frameworks consistently omit. As Ullah (2024) documents, translating governance adoption into control effectiveness requires explicit linkage to enforceable control specifications. Saleh and Abdulsalam (2025) demonstrate that governance conformance depends on mapping policy requirements to evidence types that audit processes can verify. This assignment directly implements the operationalization mechanism that Rahman et al. (2026) and the broader literature reviewed in Section 2 do not provide.

The fifth column maps each control to one of seven AI lifecycle phases: Strategy and Design, Data Preparation, Development, Procurement, Deployment, Operations, or Retirement. Nicho (2018) shows that lifecycle-process integration moves governance beyond compliance documentation toward operational enforcement. The sixth column specifies Governance Artifacts, the documentary outputs that demonstrate control implementation and provide the evidentiary basis for audit. Hevner et al. (2004) show that designed artifacts must produce tangible, evaluable outputs. Katakam (2025) shows that governance controls require stage-specific enforcement mechanisms with concrete outputs. Together, the six columns form a continuous traceability chain from governance intent to implementation evidence.

3.3 The Sixteen Governance Controls: Design Logic and Mapping Rationale

Having established the six-column alignment pathway, this section documents the design logic and mapping rationale for the 16 governance controls that populate it. The LIAGCF comprises 16 controls, evenly distributed across the four NIST AI RMF core functions: GOVERN, MAP, MEASURE, and MANAGE. Four controls per function provide sufficient granularity to demonstrate structural logic and internal alignment across control categories, SP 800-53 families, and lifecycle phases, within the bounded scope appropriate for a DSR artifact under expert panel evaluation. The selection reflects the scope of each AI RMF function as defined in NIST AI RMF 1.0 (NIST, 2023) and the precursor research that produced the initial control architecture.

GOVERN Function: Four Administrative Controls

The GOVERN function establishes the policies, processes, roles, and accountability structures that enable AI risk management. All four GOVERN controls are Administrative. A-1 and A-2 are anchored to Strategy and Design, as governance structures must precede system acquisition, design, or deployment. A-3 (third-party AI risk governance) is anchored to Procurement, and A-4 (AI audit and governance review cadence) is anchored to Operations. Administrative controls, defined as policies, governance structures, audit frameworks, and organizational roles that establish governance authority (Hong et al., 2003), are appropriate for function-level obligations. Willie (2025) establishes that effective AI governance requires integrating human supervision and policy protocols at the administrative level. Saleh and Abdulsalam (2025) confirm that administrative requirements must be specified before operational controls can be effectively deployed. SP 800-53 families assigned include Planning (PL), Program Management (PM), Supply Chain Risk Management (SR), System and Services Acquisition (SA), and Assessment, Authorization and Monitoring (CA). The SR assignment addresses a gap Rahman et al. (2026) acknowledge but do not resolve within an enterprise GRC architecture.

MAP Function: Two Administrative and Two Technical Controls

The MAP function directs organizations to identify, categorize, and analyze AI-specific risks. The LIAGCF assigns two Administrative and two Technical controls to MAP, reflecting the function's dual nature: establishing administrative risk-identification processes and implementing technical mechanisms to characterize system-level risks. Hundeyin et al. (2025) show that privacy-preserving AI governance maps administrative and technical controls to distinct categories at risk-identification stages. Katakam (2025) reports that governance gates and metric gates, functionally equivalent to the Administrative and Technical categories, operate at the same lifecycle point. The two Administrative MAP controls are anchored to Development (A-5, requiring AI risk assessments prior to deployment) and Strategy and Design (A-6, categorizing AI systems by risk level). These controls assign RA and CA families to A-5 and RA and PL families to A-6 for risk assessment and AI system risk-level categorization. The two Technical MAP controls span Development and Data Preparation. T-1 (bias and disparate impact assessment) is anchored to Development with RA and SI families. T-2 (data provenance and lineage) is anchored to Data Preparation with SA and SR families, addressing model parameter security, data integrity, and adversarial input controls (Sudarsan et al., 2025).

MEASURE Function: Three Technical and One Operational Control

The MEASURE function directs organizations to analyze and assess AI risk. Three of the four controls are Technical, reflecting the measurement-intensive nature of model evaluation, bias assessment, and performance monitoring. One is Operational, addressing procedural governance of the assessment process. Technical controls are appropriate because MEASURE outputs, including bias metrics, performance benchmarks, and drift detection results, are generated through technical mechanisms (Reuel & Undheim, 2024). SP 800-53 families for Technical MEASURE controls include Assessment, Authorization and Monitoring (CA), System and Services Acquisition (SA), System and Information Integrity (SI), and Audit and Accountability (AU). The single Operational control, anchored to Operations, assigns CA to address ongoing compliance validation obligations. Katakam (2025) demonstrates that governance reviews must be embedded as operational procedures at each lifecycle stage, not treated as one-time assessments.

MANAGE Function: Four Operational Controls

The MANAGE function directs organizations to prioritize and respond to AI risks through operational mechanisms. All four MANAGE controls are Operational, consistent with the function's emphasis on runtime enforcement, incident response, continuous monitoring, and lifecycle management. Operational controls, defined as runtime monitoring protocols, incident

response procedures, anomaly detection, and integration of human oversight, are appropriate for post-deployment governance activities (Katakam, 2025). The four controls span Data Preparation/Operations, Operations, and Retirement. Data Preparation/Operations and Operations controls assign the Incident Response (IR), Configuration Management (CM), Access Control (AC), and Identification and Authentication (IA) families. The Retirement control assigns the Media Protection (MP), System and Services Acquisition (SA), and CM families, addressing secure model disposal, data deletion, and access revocation. This directly addresses the gap Frimpong (2026) documents: retired AI systems generate residual risks that existing frameworks do not account for. Khalid and Jaffery (2025) document the end-of-life obligations this control fulfills. Loucif et al. (2025) confirm that only a small minority of lifecycle governance frameworks include Retirement as a formally specified phase with control assignments, a structural gap this study fills through the LIAGCF.

3.4 The 16-Control Sample: Scope Justification

The 16 controls represent a purposive cross-sectional sample that demonstrates structural logic, not an exhaustive governance catalog. Distribution ensures that every relevant combination of AI RMF function, control category, and lifecycle phase is represented: all four functions have four controls; all three categories are covered; all seven phases are addressed; and multiple SP 800-53 families are instantiated. This coverage is sufficient to demonstrate that the alignment pathway functions as designed and to support expert evaluation of internal coherence and structural logic.

The scope is consistent with DSR methodology. Hevner et al. (2004) establish that DSR artifacts must demonstrate utility through evaluable outputs, not exhaustive enumeration. Peffers et al. (2007) define DSR evaluation as assessing how well the artifact supports a solution, a determination enabled by cross-sectional coverage. Venable et al. (2016) confirm that naturalistic expert evaluation is appropriate at this validation stage. Gupta (2025) documents that existing frameworks provide no systematic approach to embedding governance at scale; demonstrating the architecture across a bounded control set is the appropriate first validation step.

Future research extending the LIAGCF to a comprehensive catalog covering all applicable SP 800-53 families across all lifecycle phases would constitute a necessary subsequent design cycle iteration. The present study's contribution is establishing the structural architecture and providing expert-validated evidence that it is sound.

3.5 Framework Alignment Logic

Section 3.4 defines the scope of the 16-control sample; this section explains how the LIAGCF's three structural layers integrate to form a unified governance architecture. The LIAGCF is designed to deliver its governance function through three structural layers: alignment with the NIST AI RMF core functions, operationalization via SP 800-53 control family assignments, and anchoring in specific AI lifecycle phases. These layers are mutually constitutive. Each is necessary for the others to enable verifiable governance; no single layer is sufficient on its own.

The first layer, AI RMF functional alignment, situates the LIAGCF within the risk management vocabulary established by NIST AI RMF 1.0 for federal and enterprise AI governance (NIST, 2023). Alignment with GOVERN, MAP, MEASURE, and MANAGE ensures that each control traces to an organizational risk management obligation. As Gupta (2025) documents, AI RMF alignment at the functional level produces principles rather than actionable requirements without a second operationalization layer. Paul Greene (2025) identifies this structural problem and proposes mapping AI RMF functions to the control specifications that cybersecurity governance programs already use, a precise argument this study instantiates through the LIAGCF.

The second layer, SP 800-53 operationalization, bridges the principles-to-practice gap by assigning each control to a specific SP 800-53 control family. This translates governance

intent into an auditable control specification that enterprise GRC programs can implement, test, and report against. Ullah (2024) documents that translating governance adoption into control effectiveness requires an explicit linkage between intent and enforceable specifications. Without lifecycle anchoring, SP 800-53 assignments remain context-free and practically unenforceable at any specific lifecycle point.

The third layer, lifecycle phase anchoring, specifies which of the seven AI lifecycle phases each control applies to. Gbabo et al. (2022) establish that lifecycle-based integration makes security controls contextually actionable. Nicho (2018) demonstrates that governance frameworks require lifecycle-process integration to function as enforcement mechanisms. Rashid et al. (2026) confirm that lifecycle-integrated NIST-aligned architectures can simultaneously enforce governance and operational security obligations. Together, the three layers constitute the unified governance architecture that the preceding literature review establishes is missing, and that this study is designed to supply through the LIAGCF.

4. Research Methodology

4.1 Research Paradigm: Design Science Research

Design Science Research is the appropriate methodology for this study because the research objective is to produce, evaluate, and communicate a purposefully designed governance artifact that addresses a defined organizational problem (Hevner et al., 2004). The artifact produced is the LIAGCF, a governance control framework designed to address the structural disconnect between NIST AI RMF guidance and enterprise cybersecurity GRC implementation, as documented in the preceding literature review.

The LIAGCF's design and evaluation satisfy all seven DSR guidelines established by Hevner et al. (2004), confirming the methodological legitimacy of the study's approach. The LIAGCF constitutes a viable artifact that addresses a demonstrable problem in operationalizing AI governance (Guideline 1: Design as an Artifact). The problem, the absence of an operationalized, control-level, lifecycle-integrated AI governance structure for enterprise GRC programs, is documented across sources reviewed in Section 2 (Guideline 2: Problem Relevance). The artifact is evaluated through structured expert review, as described in Section 4.2 (Guideline 3: Design Evaluation). The study contributes a framework architecture that integrates NIST AI RMF functions with SP 800-53 control families across a seven-phase AI lifecycle. It applies administrative, technical, and operational control categorization to AI governance and provides a validated expert evaluation instrument (Guideline 4: Research Contributions). The study is grounded in established control theory, information security governance literature, and DSR methodology (Guideline 5: Research Rigor). Artifact design followed a systematic, iterative process structured around the six-column alignment pathway (Guideline 6: Design as a Search Process). Results are communicated to academic and practitioner audiences through the peer-reviewed publication process (Guideline 7: Communication of Research).

The six-stage DSR process model established by Peffers et al. (2007) operationalizes these guidelines within the study's design sequence. The six stages are problem identification and motivation, definition of solution objectives, design and development of the LIAGCF, demonstration through structured expert review, evaluation of expert responses against the five research questions, and communication through this paper. This model has been adopted in recent cybersecurity governance studies, including Melaku (2023) and Wijaya et al. (2026). The DSR evaluation phase is guided by the FEDS framework (Venable et al., 2016), which identifies structured expert review as an appropriate evaluation strategy for governance artifacts at the formative validation stage.

4.2 Evaluation Strategy: Structured Expert Review

Structured expert review is implemented in this study through a purposively sampled panel of credentialed practitioners, evaluated against five dimensions aligned with the five research questions. FEDS distinguishes between formative and summative evaluation and between artificial and naturalistic settings, providing explicit guidance on which approach is appropriate at each development stage (Venable et al., 2016). At the formative validation stage, FEDS recommends naturalistic expert evaluation. Venable et al. (2012) establish that expert panels are appropriate for DSR artifact assessment at early validation stages, where the evaluative question concerns conceptual soundness and alignment with domain requirements rather than operational effectiveness.

Structured expert review is preferred over alternative evaluation approaches because the LIAGCF's face validity requires judgments from domain practitioners with direct experience in the relevant governance environments. Hevner et al. (2004) establish that DSR evaluation must demonstrate that the artifact provides a solution to the identified problem; structured expert review supports this by systematically eliciting practitioner judgments on the dimensions the framework is intended to address. Motiang and Mamorobela (2025) evaluate an information security governance framework through a semi-structured expert review with ten specialists. Rathnasinghe et al. (2026) document purposive sampling, thematic saturation assessment, and analytical trail documentation as the methodological standard for governance framework expert evaluation. Expert review is operationalized through a structured written instrument administered via Alchemer, as described in Section 4.3.

4.3 Survey Instrument Design

The survey instrument is structured into three functional segments to ensure that eligibility verification precedes artifact evaluation and that all substantive responses reflect direct engagement with the LIAGCF. Questions 1 through 4 constitute eligibility screening: participants confirm at least five years of professional experience in cybersecurity, risk management, compliance, or AI governance; affirm familiarity with the NIST framework; and self-identify their domain expertise and organizational sector. Question 5 is a confirmation gate that requires participants to verify they reviewed the LIAGCF six-column table before proceeding. Responses indicating they did not review are treated as incomplete and excluded.

The five substantive questions (Questions 6 through 10) constitute the evaluation core of the instrument, each mapped directly to a research question as defined in Section 1.3. Question 6 maps to RQ1 and addresses the structural alignment of the LIAGCF's 16 governance controls with the four NIST AI RMF core functions: GOVERN, MAP, MEASURE, and MANAGE. Question 7 maps to RQ2 and assesses whether SP 800-53 control family assignments appropriately operationalize each governance objective. Question 8 maps to RQ3 and addresses the comprehensiveness of lifecycle coverage across the seven AI system phases, from Strategy and Design through Retirement, and identifies governance domains or risk considerations that remain insufficiently addressed. Question 9 maps to RQ4 and examines whether the LIAGCF's integration of NIST AI RMF functions and SP 800-53 control families within a unified, lifecycle-aligned structure represents a meaningful and distinctive contribution to AI governance practice and scholarship. Question 10 maps to RQ5 and elicits expert assessment of the operational barriers organizations would encounter when integrating the LIAGCF into existing GRC infrastructure. The instrument also includes Question 11, an optional open-ended field for any additional observations (reported in Section 5.7), and Question 12, a follow-up consent question (reported in Appendix C, Table C4).

All five substantive questions are open-ended. Likert-scale formats are not used because the LIAGCF is a novel, multi-component artifact evaluated for conceptual soundness rather than for user acceptance or implementation performance. Open-ended questions allow experts to articulate their reasoning, identify control assignments requiring revision, and surface

dimensions that a fixed-response format would not capture (Braun & Clarke, 2006). Rahayu et al. (2025) and Ramli and Darus (2026) similarly employ discursive expert assessment for governance artifact evaluation. The FEDS framework (Venable et al., 2016) supports this approach for evaluative questions concerning conceptual soundness and domain fit. Responses are analyzed using Braun and Clarke's (2006) six-phase thematic analysis process, as described in Section 4.5.

4.4 Expert Qualification Criteria and Sampling

Expert participants are recruited through purposive sampling, in which selection criteria are based on the evaluative capacity needed to assess the artifact rather than on statistical representativeness. Purposive sampling is methodologically necessary here. The evaluative questions concern domain-specific judgments about structural logic, control category alignment, lifecycle sequencing, and GRC applicability, all of which require professional expertise in cybersecurity, risk management, compliance, and NIST frameworks (Rathnasinghe et al., 2026).

Eligibility criteria require at least five years of professional experience in cybersecurity, risk management, compliance, or AI governance, along with a working familiarity with NIST frameworks sufficient to evaluate SP 800-53 control family assignments and NIST AI RMF functional mapping. These criteria align with selection standards in published DSR governance studies, where expert qualification is defined by domain-specific security or governance expertise (Motiang & Mamorobela, 2025).

A panel of 10 complete responses is appropriate for DSR artifact evaluation at the formative validation stage. Peffers et al. (2007) support expert review without prescribing minimum panel sizes. Venable et al. (2016) establish that small expert panels constitute appropriate naturalistic evaluation at early development stages; Venable et al. (2012) confirm this for formative DSR assessment. Tremblay et al. (2010) document that panel composition and domain expertise, not numerical size, determine evaluative quality. Motiang and Mamorobela (2025) and Rathnasinghe et al. (2026) both employ ten-expert panels for information security governance artifact evaluation, providing direct precedent. The combined effect of purposive sampling, domain-specific eligibility criteria, and a panel size consistent with FEDS guidance is to ensure that the expert evaluation produces face and construct validity evidence appropriate to the LIAGCF's formative validation stage.

4.5 Unit of Analysis

The unit of analysis is the expert respondent as a source of evaluative judgment about the LIAGCF. Findings are organized thematically across the five evaluation dimensions corresponding to RQ1 through RQ5. Individual responses are not attributed or cross-tabulated; instead, they are coded and analyzed for thematic patterns across the panel as a whole.

4.6 Data Analysis, Ethical Considerations, and Positionality

Thematic analysis following Braun and Clarke's (2006) six-phase process was used to analyze the expert responses collected through the instrument described in Section 4.3. In Phase 1, all responses were read multiple times to develop broad familiarity with the dataset before coding. In Phase 2, initial codes were generated through open coding of all meaningful data segments. In Phase 3, codes were grouped into candidate themes organized around the five evaluation dimensions. In Phase 4, candidate themes were reviewed against the full coded dataset to confirm internal coherence and distinctiveness. In Phase 5, each theme was defined analytically and named to reflect its substantive content. In Phase 6, findings are reported by evaluation dimension, with representative anonymized quotations illustrating each dominant theme.

Participation was entirely voluntary. All responses are kept strictly confidential and reported anonymously. No identifying information is linked to any individual response. The study was conducted in accordance with ethical research principles governing voluntary

participation and informed consent. Questions may be directed to stanfield.miranda@gmail.com.

A distinct methodological consideration concerns the researcher's positionality: the researcher who developed the LIAGCF is also the primary analyst of expert responses evaluating it. This dual role introduces the potential for confirmation bias. To address this risk, the researcher committed to four reflexivity procedures prior to analysis. These were coding all responses in full before drawing thematic conclusions, actively seeking and prominently reporting disconfirming evidence, preserving the full codebook as an audit trail (Appendix D, Table D1), and reporting framework revision recommendations from expert criticism in Section 6.6 with the same analytical weight as validating findings. These procedures mitigate but do not eliminate positionality risk; an independent external audit of the thematic analysis would further strengthen confirmability.

4.7 Validity and Reliability Considerations

The procedural decisions documented in Sections 4.1 through 4.6 are evaluated using the four-criterion trustworthiness framework established by Lincoln and Guba (1985): credibility, transferability, dependability, and confirmability.

Credibility is established through three procedural mechanisms. Eligibility screening ensures that all participants have at least five years of domain-relevant experience and familiarity with the NIST framework. The confirmation gate ensures that all substantive answers reflect direct engagement with the artifact. Thematic analysis, following Braun and Clarke's (2006) six-phase process, is applied consistently, ensuring that analytical conclusions are derived from the data rather than from researcher presupposition.

Transferability is supported by the LIAGCF's explicit design for enterprise cybersecurity GRC programs across federal and commercial contexts. The purposive sample includes participants from diverse professional domains and organizational settings, consistent with Rathnasinghe et al.'s (2026) practice of sampling for governance perspective diversity. The study does not claim transferability beyond the scope defined by eligibility conditions; sufficient methodological description is provided for readers to assess contextual applicability (Lincoln & Guba, 1985).

Dependability is supported by a standardized written instrument administered via Alchemer, ensuring that all participants encounter identical questions in the same sequence. The instrument's design, RQ-to-question mapping, and analysis protocol are documented in sufficient detail to permit replication (Venable et al., 2016).

Confirmability is established through the documentary audit trail maintained throughout the study. Instrument questions, raw expert responses, thematic coding records, and the mapping of coded themes to reported findings are preserved and available for external audit. This approach aligns with Rahayu et al.'s (2025) iterative analysis with documented internal validation and with Lincoln and Guba's (1985) requirement for an explicit, traceable record connecting raw data to analytical conclusions.

5. Findings

5.1 Expert Panel Profile

The final analytic sample of N = 10 comprises credentialed practitioners who met all eligibility and confirmation gate criteria for valid expert evaluation of the LIAGCF. Eligibility criteria required a minimum of five years of professional experience in cybersecurity, risk management, compliance, or AI governance; working familiarity with NIST-based governance frameworks; and confirmed review of the LIAGCF governance control mapping table prior to responding. Three additional PDF exports left the Q5 confirmation gate field blank due to a known Alchemer rendering behavior on certain browser and mobile submissions. Raw platform data confirmed that

all three respondents met the confirmation gate criterion, as evidenced by their progression to subsequent survey items. Those responses were excluded in accordance with the instrument's specified inclusion criteria.

Panel members' professional backgrounds span the experience levels and framework familiarity distributions relevant to the LIAGCF's evaluation dimensions. Professional experience was distributed across three levels: five respondents (50%) reported 5 to 10 years, three (30%) reported 11 to 15 years, and two (20%) reported 16 to 20 years. SP 800-53 familiarity was evenly distributed: five respondents (50%) reported extensive implementation or assessment experience, and five (50%) reported working familiarity. AI RMF familiarity showed greater variance: one respondent (10%) reported extensive experience, six (60%) reported working familiarity, two (20%) reported limited familiarity, and one (10%) reported no prior familiarity. The eligibility criterion specified broad working familiarity with NIST-based governance frameworks; AI RMF familiarity was not a standalone eligibility requirement, and all ten respondents satisfied the stated eligibility criteria. This distribution is acknowledged as a limitation in Section 6.7.

The panel's primary areas of expertise align directly with the governance domains the LIAGCF is designed to address, supporting the face validity of the expert evaluation. Eight respondents (80%) identified Cybersecurity Governance; seven each (70%) identified Compliance and Audit and Federal GRC Programs; six (60%) identified Enterprise Risk Management; four (40%) identified Commercial GRC Programs; and three (30%) identified AI/ML Security. The panel's collective expertise profile is consistent with the enterprise and federal GRC practitioner audience this study addresses through the LIAGCF, confirming that the evaluative judgments reported in Sections 5.2 through 5.6 are grounded in directly relevant domain expertise. Table 1 presents the complete demographic profile.

Table 1. Expert Panel Demographic Profile (N = 10)

Characteristic	Category	n (%)
Years of Professional Experience	5–10 years	5 (50%)
	11–15 years	3 (30%)
	16–20 years	2 (20%)
SP 800-53 Familiarity	Extensive experience	5 (50%)
	Working familiarity	5 (50%)
AI RMF Familiarity	Extensive experience	1 (10%)
	Working familiarity	6 (60%)
	Limited familiarity	2 (20%)
	No familiarity	1 (10%)
Primary Expertise Area (select all that apply)	Cybersecurity Governance	8 (80%)
	Compliance & Audit	7 (70%)
	Federal GRC Programs	7 (70%)
	Enterprise Risk Management	6 (60%)
	Commercial GRC Program	4 (40%)
	AI/ML Security	3 (30%)

Note. SP 800-53 = NIST Special Publication 800-53. AI RMF = NIST Artificial Intelligence Risk Management Framework. Expertise areas reflect multiple-selection responses; percentages sum to more than 100%.

5.2 RQ1 Findings: NIST AI RMF Structural Alignment

Expert assessment of RQ1 provided strong validation of the LIAGCF's structural alignment with the NIST AI RMF core functions, with nine of ten respondents confirming that the control distribution reflects a logically sequenced risk management lifecycle. Thematic analysis of Q6 responses yielded three themes: structural alignment broadly validated (T1-A), cross-functional control behavior at function boundaries (T1-B), and the need for clarification of the MANAGE/MEASURE boundary (T1-C).

Theme T1-A, Structural Alignment Broadly Validated, was the dominant finding, supported by nine of ten respondents. Expert respondents confirmed that the control distribution across GOVERN, MAP, MEASURE, and MANAGE reflects a logically sequenced risk management lifecycle consistent with the NIST AI RMF 1.0 functional intent. GOVERN is structured to anchor organizational policy before system acquisition; MAP is structured to identify and categorize AI-specific risks within organizational settings. MEASURE is designed to analyze and assess AI risks through technical evaluation and monitoring, and MANAGE is structured to enforce governance obligations during deployment and retirement. One respondent characterized the progression as "mirroring the risk management lifecycle articulated by NIST and demonstrating a coherent governance logic that integrates policy formation, risk identification, measurement, and mitigation." Another respondent assessed one-to-one alignment and identified no misalignment.

Theme T1-B, Cross-Functional Control Behavior at Function Boundaries, was identified by six respondents. Certain controls, particularly the bias and discriminatory impact assessment control and the data provenance control, exhibit functional characteristics of both MAP and MEASURE. The bias assessment control, assigned to MAP for risk identification, also involves quantitative model testing, more characteristic of MEASURE. Respondents characterized this as inherent to integrated governance frameworks rather than a structural flaw. One respondent noted that "these overlaps do not undermine the framework but suggest an opportunity to explicitly acknowledge cross-functional control behavior in future iterations."

Theme T1-C, MANAGE/MEASURE Boundary Clarification Needed, was identified by two respondents. These respondents noted that audit trail and logging controls assigned to MEASURE exhibit accountability characteristics more typically associated with GOVERN. One respondent observed that "MANAGE can become a catch-all for controls that didn't fit neatly elsewhere, particularly incident response, model retraining triggers, and decommissioning protocols," cautioning that this risks underspecifying which controls are truly reactive versus those requiring ongoing monitoring. This finding is characterized as a technical refinement consideration for future iterations rather than a structural deficiency.

5.3 RQ2 Findings: SP 800-53 Control Family Operationalization

Expert assessment of RQ2 confirmed that the SP 800-53 control family assignments provide a defensible operational foundation and identified targeted expansion opportunities for AI-specific risk dimensions. Thematic analysis of Q7 responses yielded three themes: operationalization broadly validated (T2-A), targeted expansion recommendations (T2-B), and SP 800-53 AI-native tension (T2-C).

Theme T2-A, SP 800-53 Operationalization Broadly Validated, was the dominant finding, supported by seven of ten respondents. Expert respondents confirmed that the control family assignments provide a defensible operational foundation for each governance objective. Well-matched families included: PM and PL for governance policy and accountability; RA for risk assessment and system categorization; CA for validation and review; AU for logging and traceability; IR for AI incident response; and CM for model integrity protection. One respondent stated that "the mapping to SP 800-53 control families

significantly improves integration potential because many organizations already structure their governance processes around these control families."

Theme T2-B, Targeted Expansion Recommendations, was identified by seven respondents. The most consistent recommendation, raised by two respondents, was to add the Personally Identifiable Information Processing and Transparency (PT) family to the bias and discriminatory impact assessment control. One respondent explained that "PT seems like a natural fit here given that affected populations are often defined by personal data used in training or inferencing" and that the addition "would strengthen that row considerably." Another respondent recommended adding AU to the Retirement phase control, noting that "demonstrating compliant decommissioning typically requires documented evidence" and that AU would "add important coverage" in an audit context. Additional recommendations included CA for fairness evaluation validation (three respondents) and SI for training data integrity governance (two respondents). These recommendations characterize expansion opportunities rather than corrections.

Theme T2-C, SP 800-53 AI-Native Tension, was identified by one respondent, who recommended "a hybrid mapping approach, retaining SP 800-53 families for compliance continuity while layering AI-specific controls from 600-1 and the AI RMF Playbook onto rows where the fit is weakest." This observation may inform future development but is outside the scope of the current study.

5.4 RQ3 Findings: AI Lifecycle Representation

Expert assessment of RQ3 validated the LIAGCF's seven-phase lifecycle structure as appropriately comprehensive and identified transparency, explainability, and post-deployment fairness monitoring as the most significant governance gaps. Thematic analysis of Q8 responses produced four themes: lifecycle coverage comprehensive (T3-A); transparency, explainability, and post-deployment fairness monitoring as priority gaps (T3-B); deployment-phase governance underrepresented (T3-C); and retraining governance undertheorized (T3-D).

Theme T3-A, Lifecycle Coverage Comprehensive, was the dominant finding, supported by seven of ten respondents. Expert respondents validated the seven-phase structure as appropriately comprehensive. Data Preparation was recognized by two respondents as addressing a stage that governance frameworks frequently overlook. One respondent stated that "including Data Preparation as its own phase is a good decision because a lot of governance frameworks gloss over that stage even though data quality and provenance issues affect nearly everything downstream." Procurement was noted by two respondents as important and often absent from comparable frameworks, particularly relevant when organizations consume third-party AI. Two respondents acknowledged Retirement as valuable; one noted its consistent absence from similar work.

Theme T3-B, Transparency, Explainability, and Post-Deployment Fairness Monitoring as Priority Gaps, was identified by five respondents and is the most prevalent critical finding. Expert respondents consistently identified the lack of explicit governance controls for model transparency, explainability, and ongoing fairness monitoring in the Operations phase as the most significant gap. One respondent noted that "there is no dedicated control for transparency or explainability in the post-deployment phase, which is becoming harder to ignore given where AI regulation is heading globally." Another observed that "bias and disparate impact do not stop being relevant once a model is deployed" and that adding a fairness monitoring control in Operations "would better reflect how these risks actually behave over time." This theme is addressed in Section 6.6 as a priority revision.

Themes T3-C and T3-D, Deployment Phase and Retraining Governance, were each identified by two to three respondents. The transition within the Deployment phase could be more explicit, with controls for deployment authorization, production readiness review, and rollback criteria. One respondent noted that "in practice, deployment is one of the highest-risk

transition points for AI systems, especially when outputs begin affecting real users, business decisions, or regulated processes." Another respondent identified the absence of governance structures for retraining approval, dataset refresh, and model re-certification as a gap for continuously learning systems. These are characterized as development priorities for future iterations.

5.5 RQ4 Findings: Theoretical Contribution

Expert assessment of RQ4 yielded near-unanimous affirmation of the LIAGCF's theoretical and practical contribution, with audit traceability and translational value identified as its primary distinguishing strength. Thematic analysis of Q9 responses yielded four themes. These are consensus on meaningful contribution (T4-A), audit traceability and translational value as the primary strength (T4-B), the integrative-structural nature of the contribution (T4-C), and documentation of the control selection methodology as a critical pre-publication revision (T4-D).

Theme T4-A, Consensus on Meaningful Contribution, was the dominant finding, supported by nine of ten respondents. Expert respondents uniformly affirmed that the LIAGCF represents a meaningful contribution. Multiple respondents framed the contribution as the study's resolution of an underaddressed structural gap between principled AI risk management and operationalized cybersecurity enforcement. Two respondents expressed the view that the framework has sufficient structural soundness for operational piloting. One respondent described the framework as providing "a traceable path from governance intent all the way through to an audit artifact," characterizing this as a practically valuable contribution, "not something that has been structured quite this way before."

Theme T4-B, Audit Traceability and Translational Value as Primary Strength, was identified by eight of ten respondents. Respondents consistently highlighted the ability to translate AI governance concepts into auditable, enforceable controls that are traceable from governance objectives through lifecycle phases to implementation artifacts as the core practical value distinguishing the LIAGCF from existing frameworks. One respondent described this as "operational traceability: if each AI risk objective maps to measurable 800-53 controls across lifecycle phases, that creates audit defensibility and implementation clarity." Another characterized the framework's value as "bridging two worlds that are often discussed separately," noting that it "translates AI governance concepts into structures that enterprise GRC teams already use," enabling GRC teams, auditors, and risk owners to engage with AI governance through familiar structures.

Theme T4-C, Integrative-Structural Nature of the Contribution, was identified by two respondents. They characterized the LIAGCF's contribution as integrative and structural rather than conceptually original. One respondent stated that "the individual components, RMF alignment, 800-53 mapping, lifecycle framing, each exist in the literature independently" and that "the contribution is integrative and structural rather than conceptually original, which is a legitimate and valuable form of scholarly contribution but should be framed as synthesis and operationalization rather than theoretical invention to withstand peer review scrutiny." This assessment aligns directly with the contribution framing established in Section 1.5.

Theme T4-D, Methodology Documentation as Critical Pre-Publication Revision, was identified by four respondents as the most consistent critical finding. Respondents called for explicit documentation of the selection of the 16 controls, the assignment of SP 800-53 families, and the criteria for anchoring lifecycle phases. One respondent stated that "the framework presents its alignment decisions without explaining the process behind them" and that "this will be a question evaluators ask." The methodological justification is addressed in Section 3.3 and Appendix E (Table E1); taken together, these elements constitute the documentation respondents identified as necessary for the contribution claim to withstand peer review.

5.6 RQ5 Findings: Operational Integration

Expert assessment of RQ5 confirmed that the LIAGCF is well-positioned for integration within NIST-aligned organizations, with organizational barriers identified as the primary challenge to adoption. Thematic analysis of Q10 responses produced four themes: integration is feasible for NIST-aligned organizations (T5-A), organizational barriers are the dominant adoption challenge (T5-B), technical infrastructure gaps are secondary barriers (T5-C), and maturity-tiered implementation guidance is needed (T5-D).

Theme T5-A, Integration Feasible for NIST-Aligned Organizations, was the dominant finding, supported by eight of ten respondents. SP 800-53 grounding was consistently identified as the primary enabler for organizations whose control environments are already structured around SP 800-53 families. One respondent with extensive federal GRC experience noted that "in mature federal agencies or contractors that follow established federal security rules, most of the framework's controls are already in place or build directly on existing requirements" and that "the framework's lifecycle approach fits well with the ongoing, risk-based authorization process used in federal security programs, allowing AI-specific adjustments without adding new system silos." Another respondent noted that tying controls to lifecycle phases and familiar SP 800-53 families makes AI governance more auditable and practical to implement within existing GRC program structures.

Theme T5-B, Organizational Barriers Dominate, was cited by seven of ten respondents. Three barriers emerged consistently. Unclear AI risk ownership across legal, compliance, cybersecurity, and data science functions was identified by five respondents as the most significant challenge. One respondent noted that "without defined accountability, even a well-structured framework can become advisory rather than enforceable." Another identified a compounding structural barrier, describing the coordination demand as requiring practitioners "to simultaneously hold the AI RMF, SP 800-53, and the LIAGCF lifecycle structure in operational tension" and warning that "without explicit organizational wiring, integration stalls at the policy layer and never reaches operational practice." Incomplete AI system inventories, including AI components not captured in the organizational AI inventory, such as AI embedded in third-party vendor software, were identified by three respondents as a foundational gap. Workforce expertise gaps that combine AI governance knowledge with GRC and cybersecurity competency were identified by four respondents as a structural constraint.

Theme T5-C, Technical Infrastructure Gaps as Secondary Barriers, was identified by three respondents, who noted that several MEASURE controls assume monitoring infrastructure that many organizations have not yet built, including model drift detection tooling and AI-specific performance metrics. One respondent noted that "existing GRC platforms may not yet support AI-specific governance metrics such as model fairness or drift monitoring," which requires investment in tooling before certain controls can be operationalized.

Theme T5-D, Maturity-Tiered Implementation Guidance Needed, was identified by three respondents, who noted that the LIAGCF appears to be designed for organizations with established AI governance infrastructure. Respondents recommended tiered implementation pathways, analogous to the NIST Cybersecurity Framework's maturity tiers, to broaden applicability for less mature organizations. One respondent observed that "guidance for organizations at different maturity levels would go a long way toward making this framework more accessible and adoptable in practice."

5.7 Additional Expert Observations

Supplementary observations from Q11 yielded four actionable recommendations: a visual lifecycle diagram, explicit scope framing, clarification of the artifact column’s status, and tiered implementation guidance. One respondent recommended a visual lifecycle diagram that illustrates

the flow of governance activities across the seven phases, describing it as follows: "A simple lifecycle diagram showing how governance decisions move from system proposal to risk assessment, validation, deployment, monitoring, and retirement could make the framework easier to interpret." The same respondent suggested explicitly framing the study's scope as governance and risk management rather than AI ethics or technical model development, noting that "making that scope explicit could help readers better understand the purpose of the model." A second respondent recommended clarifying whether entries in the Governance Artifacts column are illustrative examples or required evidence specifications, stating that "that clarification would help both academic reviewers and practitioners understand how to interpret the table."

One respondent raised a structural observation on the equal distribution of controls, noting that "having exactly four controls per function looks clean, but it may not reflect how governance work is actually distributed across the lifecycle" and that "MANAGE, for example, tends to carry more operational weight than GOVERN in day-to-day practice." Another respondent recommended making the implicit assumption of organizational maturity explicit, proposing that "a tiered implementation pathway, analogous to the NIST Cybersecurity Framework's maturity tiers, would significantly broaden the framework's applicability and impact without compromising its rigor for sophisticated adopters."

6. Discussion

6.1 Interpreting RMF Alignment Findings

The near-unanimous validation of the LIAGCF's structural alignment with the four NIST AI RMF core functions provides empirical support for this study's foundational design claim. A lifecycle-integrated governance control structure can be systematically organized around the AI RMF's functional architecture without distorting the functions' definitional intent or the controls' operational purpose. Nine of ten respondents confirmed that the control distribution reflects a logically sequenced risk management lifecycle consistent with NIST AI RMF 1.0. This directly addresses what Gupta (2025) identifies as the persistent failure of existing frameworks to translate RMF functional guidance into assignable, verifiable controls.

The cross-functional behavior identified at the MAP/MEASURE boundary is best interpreted as evidence of structural integration rather than a design error. Expert respondents framed this as inherent to integrated governance architectures; the appropriate response is to acknowledge cross-functional behavior in future iterations, not to reassign. The minority MANAGE/MEASURE finding reflects a genuine architectural tension: audit logging controls serve both measurement and accountability purposes. Future iterations may benefit from dual-function notation, consistent with the NIST Cybersecurity Framework's approach to cross-category applicability.

6.2 SP 800-53 Mapping: Validation and Refinement

Complementing the RMF alignment findings, expert validation of the SP 800-53 control family assignments confirms that the framework's operationalization layer is defensible across practitioner expertise levels. Seven of ten respondents confirmed that PM, PL, RA, CA, AU, IR, and CM are well-matched, indicating that the operationalization logic is legible across the SP 800-53 familiarity distribution.

The two highest-priority expansion recommendations, PT for bias risk governance and AU for Retirement phase decommissioning, address documented coverage gaps at the margins of an otherwise defensible operationalization structure. Adding the Personally Identifiable Information Processing and Transparency (PT) family to the bias and discriminatory impact assessment control addresses a genuine gap. Affected populations are often defined by personal data, making PT's transparency and data processing governance controls directly relevant to bias risk governance (Finch & Butt, 2025). Adding AU to the

Retirement phase control addresses the requirement for compliant decommissioning to provide traceable evidence of secure model disposal, data deletion, and access revocation. This study uses SP 800-53 as the operationalization layer by design, grounded in enterprise GRC compatibility and the catalog's broad adoption across federal and commercial programs, rather than as a claim of complete AI-specific coverage.

6.3 Lifecycle Coverage: Gaps and Opportunities

Beyond control family operationalization, expert validation of the seven-phase lifecycle structure confirms coverage of domains that comparable frameworks consistently underrepresent. The inclusion of the retirement phase is especially significant given the literature's documented absence of end-of-life lifecycle governance (Frimpong, 2026; Loucif et al., 2025).

The panel's most actionable gap is the lack of explicit post-deployment governance controls for transparency, explainability, and fairness monitoring in the Operations phase. Five respondents identified these as priority gaps, reflecting regulatory developments, including the EU AI Act and NIST AI RMF Playbook guidance, that increasingly require explicit governance for transparency and explainability. Bias and fairness are addressed in the Development phase, but these risks persist through deployment as input distributions shift and model behavior changes. This gap is prioritized as a high-priority revision in Section 6 .6. Deployment-phase and retraining governance gaps, identified by two to three respondents, are secondary findings characterized as development priorities.

6.4 Theoretical Positioning of the LIAGCF

Taken together, the RMF alignment, SP 800-53 operationalization, and lifecycle coverage findings provide the evidentiary foundation for the contribution assessment reported in RQ4. Near-unanimous expert consensus on the LIAGCF's meaningful contribution, supported by nine of ten respondents, provides face and construct validity evidence appropriate for the formative evaluation stage. The framework's value was assessed as meaningful regardless of respondents' level of AI RMF familiarity, supporting its applicability across the practitioner population this study targets. The consistent identification of audit traceability and translational value as the primary strength directly validates the paper's theoretical positioning. Multiple respondents noted that the LIAGCF is structured to provide a traceable path from governance intent to an audit artifact and to render AI governance concepts as structures that enterprise GRC teams already use. This convergent validation across federal and commercial GRC environments provides strong empirical support for the contribution claim.

Two respondents characterized the contribution as integrative and structural rather than conceptually original, noting that it synthesizes and operationalizes existing NIST standards. This validates the framing of the contribution in Section 1.5 and aligns with the DSR literature's recognition that operationalization artifacts address the gap between frameworks and operational practice, where governance programs most frequently fail. Such framing positions the LIAGCF's contribution within the DSR tradition and directly informs the practical deployment assessment in Section 6.5.

6.5 Practical Deployment Considerations

Expert assessment findings confirm that the LIAGCF's SP 800-53 grounding resolves the most common integration failure mode: adoption as a parallel governance program rather than as an extension of existing control environments. The lifecycle orientation was found to be compatible with continuous, risk-based authorization processes, enabling AI-specific governance without separate oversight structures. However, the organizational barriers identified by respondents precede framework adoption and cannot be resolved through framework design alone. Unclear AI risk ownership, incomplete AI system inventories, and workforce expertise gaps each require organizational action beyond the control architecture. The GOVERN-function controls are

designed to address accountability structures and decision rights; the MAP-function controls are designed to require AI system categorization and inventory documentation; and grounding in SP 800-53 reduces the burden of new knowledge for GRC practitioners. The LIAGCF is tool-agnostic by design, allowing organizations to operationalize controls with available capabilities and extend coverage as AI-specific tooling matures.

6.6 Framework Revisions Indicated by Expert Feedback

The gaps and barriers documented in Sections 8.1 through 8.5 are consolidated in Table 2 as a structured revision log organized by priority level. High-priority revisions address findings reported by four or more respondents or those that directly affect publication readiness. Medium-priority revisions address findings from two to three respondents that indicate meaningful governance gaps. Low-priority revisions address presentation and accessibility.

Table 2. Framework Revision Log: Expert Feedback-Indicated Revisions to the LIAGCF

Revision Area	Expert Feedback Summary	Recommended Action	Priority	Relevant Framework Section
Transparency and Explainability Governance	Five respondents identified the absence of explicit controls for model transparency, explainability documentation, and user-facing disclosure obligations as the most significant lifecycle gap.	Add a Technical control in the MAP or MEASURE function governing explainability documentation requirements; add an Operational control in MANAGE formalizing stakeholder transparency obligations in Operations phase.	**High**	Table A1; Appendix A; Section 3.3
Post-Deployment Fairness Monitoring	Five respondents noted that the fairness evaluation control is scoped to Development only, but bias and disparate impact risks persist through the full operational lifecycle.	Add a MEASURE or MANAGE operational control in the Operations phase governing ongoing fairness monitoring, threshold-based alerting, and documented remediation when disparate impact is detected post-deployment.	**High**	Control T-1; Appendix A; Section 3.3
PT Family Addition — Bias and Fairness Control	Two respondents recommended adding Personally Identifiable Information Processing and Transparency (PT) to the bias and discriminatory impact assessment control, noting that affected populations are frequently defined by personal data.	Expand SP 800-53 family assignment for Control T-1 from RA + SI to RA + SI + PT in next framework iteration.	**High**	Control T-1; Appendix E; Section 3.3
Control Selection Methodology Documentation	Four respondents identified the absence of explicit documentation for how the 16 controls were selected and how SP 800-53 families were assigned as the most critical pre-publication revision.	Section 3.3 and Appendix E of this paper provide the methodological justification for all 16 control selections and SP 800-53 family assignments. The expert validation conducted in this study constitutes the practitioner review respondents identified as necessary.	**High**	Section 3.3; Appendix E
AU Family Addition — Retirement Control	One respondent recommended adding Audit and Accountability (AU) to the Retirement phase control, noting that demonstrating compliant decommissioning requires documented evidence artifacts.	Expand SP 800-53 family assignment for Control O-5 from SA + CM + MP to SA + CM + MP + AU in next framework iteration.	**Medium**	Control O-5; Appendix E; Section 3.3
Deployment Phase Governance	Two respondents identified the Deployment phase as underrepresented, noting the	Add one Administrative or Operational control in the Deployment phase addressing	**Medium**	Appendix A; Section 3.3; Section 7.5

Revision Area	Expert Feedback Summary	Recommended Action	Priority	Relevant Framework Section
	absence of controls for deployment authorization, production readiness review, rollback criteria, and post-deployment monitoring triggers.	production readiness review and deployment authorization criteria in a future expanded control set.		
Retraining and Model Re-Certification Governance	Two respondents identified the absence of governance structures for retraining approval, dataset refresh governance, and model re-certification as a lifecycle completeness gap for systems that evolve through continuous learning.	Add a MANAGE Operational control in the Operations phase addressing retraining approval workflow, dataset refresh governance, and re-certification requirements in a future expanded iteration.	**Medium**	Appendix A; Section 7.5
Human Oversight and Override Controls	Two respondents identified human-in-the-loop oversight and decision override mechanisms as undertheorized in the current framework.	Add an Operational MANAGE control in the Deployment phase explicitly addressing human-in-the-loop review protocols and decision override procedures for high-risk AI decisions.	**Medium**	Appendix A; Section 7.4
Tiered Implementation Pathway	Three respondents noted the framework assumes organizational AI governance maturity and recommended tiered implementation pathways for organizations at different governance development stages.	Develop a maturity model overlay for the LIAGCF providing tiered implementation guidance in future research, as described in Section 7.4.	**Low**	Section 7.4; Section 7.5
Visual Lifecycle Diagram	One respondent recommended inclusion of a visual diagram illustrating how governance activities flow across the seven lifecycle phases.	Add a lifecycle flow diagram as a supplementary figure in the published version, illustrating the governance activity progression from Strategy and Design through Retirement.	**Low**	Section 5; Appendix A
Governance Artifacts Column Clarification	One respondent recommended clarifying whether governance artifacts listed in the sixth column are illustrative examples or required evidence specifications.	Add a table note to Appendix A explicitly stating that governance artifacts are illustrative examples of implementation evidence rather than prescriptive requirements.	**Low**	Appendix A; Table Note

Note. Priority levels: High = reported by four or more respondents or directly affecting publication readiness; Medium = reported by two to three respondents indicating meaningful governance gaps; Low = presentation and accessibility improvements. Abbreviations: LIAGCF = Lifecycle-Integrated AI Governance Control Framework; SP 800-53 = NIST Special Publication 800-53; AI RMF = NIST Artificial Intelligence Risk Management Framework; PT = Personally Identifiable Information Processing and Transparency; AU = Audit and Accountability; CA = Assessment, Authorization and Monitoring; SI = System and Information Integrity.

These revisions do not correct the foundational architecture, which expert respondents validated as structurally sound. They reflect the iterative refinement cycle that DSR artifact development requires after the initial expert evaluation (Hevner et al., 2004).

6.7 Limitations

Expert validation establishes face and construct validity for the LIAGCF but not operational effectiveness in deployed settings. This is methodologically appropriate for the current DSR evaluation phase; deployment testing is identified as a direction for future research.

The panel of N = 10, consistent with purposive sampling norms for DSR artifact evaluation at the formative stage (Venable et al., 2016; Motiang & Mamorobela, 2025), may be considered limited given the breadth of the enterprise AI governance practitioner community. Purposive sampling is the methodologically appropriate strategy here; no claim of generalizability to broader practitioner populations is made.

Three of ten respondents reported limited or no familiarity with the AI RMF, reflecting the current state of AI RMF adoption in enterprise GRC environments. This suggests that

some RMF alignment judgments may have been informed by SP 800-53 expertise and general governance knowledge rather than direct familiarity with the AI RMF. These responses are retained because their GRC practitioner perspective falls within the framework's intended user population. As the designer of the LIAGCF, the researcher serves a dual role as both artifact creator and main analyst, which could lead to confirmation bias. This risk was mitigated through reflexivity procedures outlined in Section 4.6. The 16 governance controls constitute a purposive cross-section intended to test structural alignment logic, not an operationally exhaustive catalog. This is a deliberate design choice appropriate for a DSR validation study, not a limitation on operational applicability. Expert qualification data was gathered through self-reported survey responses without independent verification. The substantive depth of responses provides indirect evidence of genuine domain expertise, and the anonymous, voluntary nature of participation reduces the incentive to misrepresent.

7. Conclusion

7.1 Summary of Findings

Expert assessment of the LIAGCF across five evaluation dimensions provided strong validation evidence, establishing face and construct validity for the framework as a governance artifact for enterprise cybersecurity GRC programs. This study developed and expert-validated the LIAGCF, a DSR artifact comprising 16 governance controls organized across four NIST AI RMF core functions, seven AI system lifecycle phases, and three control categories. It was evaluated through structured expert review by N = 10 credentialed practitioners.

Expert assessment produced five dominant findings. Structural alignment with the NIST AI RMF core functions was validated by nine of ten respondents. SP 800-53 control family assignments were deemed technically defensible by seven of ten respondents, with targeted expansion recommendations for PT and AU. The seven-phase lifecycle structure was deemed comprehensive by seven of ten respondents; transparency, explainability, and post-deployment fairness monitoring were identified as the most significant coverage gaps. Nine of ten respondents affirmed the LIAGCF as a meaningful scholarly and practical contribution; eight respondents identified audit traceability and translational value as its primary strengths. Integration feasibility within NIST-aligned GRC environments was confirmed by eight respondents; unclear AI risk ownership, incomplete AI system inventories, and workforce expertise gaps were identified as the primary adoption challenges. These findings provide face and construct validity for the LIAGCF.

7.2 Theoretical and Practical Contributions

The findings in Section 7.1 provide the empirical foundation for a dual contribution claim. Theoretically, the LIAGCF addresses a documented structural gap as the first lifecycle-integrated, control-categorized governance framework to operationalize both NIST AI RMF core functions and NIST SP 800-53 control families within a unified architecture, bridging the persistent disconnect between strategic AI risk oversight and operational cybersecurity enforcement. Practically, the LIAGCF equips GRC practitioners, CISOs, and federal risk management program leads with a standards-aligned, auditable governance structure that operationalizes AI risk management intent as assignable, lifecycle-anchored controls with designated implementation artifact requirements. The integrative and structural character of this contribution is analytically significant and warrants explicit framing. Its primary strength is synthesizing established NIST standards into a governance architecture that organizations can deploy, audit, and extend within existing GRC infrastructure. This operationalization gap has been documented in independent systematic reviews (Gupta, 2025; Finch & Butt, 2025) and confirmed by expert assessment. This study does not propose new governance principles; the LIAGCF is designed to provide the structural mechanism through which existing principles become enforceable.

7.3 Recommendations for Practitioners

The practical contribution documented in Section 7.2 offers five specific recommendations for GRC practitioners, CISOs, and AI governance program leads.

1. Begin with the four GOVERN-function administrative controls. These controls establish accountability structures, decision rights, and audit cadence, without which technical and operational controls cannot be meaningfully enforced. Organizations should treat these as preconditions for implementation rather than as parallel activities.
2. Map the 16 controls to the existing SP 800-53 baselines before conducting a gap analysis. Organizations under NIST-aligned frameworks already have SP 800-53 families documented in system security plans (SSPs) and risk registers. Identify which of the 13 control families are already addressed, and scope new activity to AI-specific gaps, particularly SR at procurement, AU for decision traceability, and MP and SA at retirement.
3. Treat the Retirement lifecycle phase as an immediate governance priority. Expert respondents identified decommissioning governance as a high-priority domain, consistent with the literature documenting residual AI risks from ungoverned system retirement. Organizations are advised to implement the Retirement-phase MANAGE control, covering secure model disposal, data deletion, access revocation, and audit documentation, before decommissioning occurs.
4. Before deploying technical controls, assign explicit ownership. Map each of the 16 controls to a designated responsible party: administrative controls to GRC and compliance functions, technical controls to AI engineering and data science teams, and operational controls to security operations and risk management.
5. Use the LIAGCF's six-column structure as a template for compliance evidence. The Governance Artifacts column lists the documentation that demonstrates control implementation. Organizations subject to AI-related audit or regulatory review can use these specifications to structure evidence collection before assessment.

7.4 Recommendations for Framework Refinement

Distinct from practitioner implementation guidance, expert feedback from the assessment study has identified three priority areas for refinement in future iterations of the LIAGCF.

First, explicit governance controls for transparency, explainability, and human oversight should be implemented. A future iteration should include at least one Technical control within MAP or MEASURE that addresses explainability documentation, and at least one Operational control within MANAGE that formalizes human oversight and override procedures for high-stakes AI decisions.

Second, SP 800-53 family assignments should be expanded to include the bias and fairness assessment control and the model validation control. Expert respondents recommended adding CA for fairness evaluation validation, SI for model robustness testing, and PT for controls governing AI systems that process personally identifiable information.

Third, a maturity model overlay should be developed to enable organizations to assess their governance posture and identify prioritized implementation pathways. This would allow organizations to begin at a foundational level using the GOVERN administrative controls and advance progressively as governance capacity develops.

7.5 Future Research Directions

Beyond framework refinement, the present study outlines a research agenda spanning six interconnected directions.

The most immediate priority is empirical validation of deployment. Face and construct validity have been established through expert assessment; the necessary next stage is

naturalistic deployment in enterprise or federal GRC programs, measuring control implementation outcomes against predefined governance effectiveness criteria.

A second direction is quantitative validation through a larger expert-panel study. Such a study, administering a revised instrument with Likert-scale items mapped to the five research questions to 50 or more respondents would enable statistical analysis of alignment ratings, inter-rater reliability, and comparative analysis across expertise profiles.

Third, sector-specific adaptation offers both practical and scholarly opportunities. Developing validated LIAGCF variants for healthcare, financial services, and national security contexts would extend the framework's practical reach and yield sector-specific theoretical contributions.

Fourth, standards integration research would clarify the LIAGCF's positioning within the evolving global AI governance landscape. Examining alignment with ISO/IEC 42001:2023 and EU AI Act compliance obligations would broaden its relevance from a NIST-centric architecture to a multi-standard enterprise governance instrument.

Fifth, longitudinal research on governance effectiveness represents the ultimate empirical test. Tracking control effectiveness across the full system lifecycle, measured at defined intervals from Strategy and Design through Retirement, would assess whether lifecycle-anchored, control-categorized governance yields measurably better AI risk outcomes than existing approaches. This research program would provide the definitive empirical test of the LIAGCF's central design claim: that lifecycle-integrated, control-categorized governance yields measurably better AI risk outcomes than existing approaches.

Sixth, a tiered implementation pathway should be developed as described in the framework refinement recommendations in Section 7.4, extending the maturity model overlay concept into a validated research instrument suitable for empirical evaluation across organizations at different governance development stages.

References

Adebayo, A. (2023). The SAIS-GRC framework: Engineering trust and secure, agile systems for proactive AI governance and compliance. *Iconic Research and Engineering Journals, 7*(5). https://doi.org/10.64388/irev7i5-1713349

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology, 3*(2), 77–101. https://doi.org/10.1191/1478088706qp063oa

Butt, T., Iqbal, M., & Arshad, N. (2025). From policy to pipeline: A governance framework for AI development and operations pipelines. *IEEE Access, 14,* 1373–1397. https://doi.org/10.1109/access.2025.3647479

Casimir, P. (2026). The governance of intelligence: Scaling trusted data through machine learning, artificial intelligence, and large language models. *Journal of Artificial Intelligence & Cloud Computing, 5*(1), 1. https://doi.org/10.47363/JAICC/2026(5)510

Catapang, J. K. (2026). Building the ethical AI framework of the future: From philosophy to practice. *AI and Ethics*. Advance online publication. https://doi.org/10.1007/s43681-025-00683-6

Egwuatu, O. (2025). Ethical and governance challenges of AI in information systems: Toward responsible adoption in enterprise systems. *World Journal of Advanced Research and Reviews, 27*(2), 1744–1751. https://doi.org/10.30574/wjarr.2025.27.2.3064

Essien, I. A., Cadet, E., Ajayi, J. O., Erigh, E. D., & Obuse, E. (2025). AI-driven continuous compliance and threat intelligence model for adaptive GRC in complex digital ecosystems. *Computer Science & IT Research Journal, 6*(7), 403–417. https://www.researchgate.net/publication/394998384

Faruq, M. O. (2025). A meta-analysis of cybersecurity framework integration in GRC platforms: Evidence from U.S. enterprise audits. *Journal of Sustainable Development and Policy, 1*(1), 224–249. https://doi.org/10.63125/kwhkmb57

Finch, W. W., & Butt, M. (2025). Gaps in AI-compliant complementary governance frameworks' suitability (for low-capacity actors), and structural asymmetries in the compliance ecosystem: A systematic review. *Journal of Cybersecurity and Privacy, 5*(4), 101. https://doi.org/10.3390/jcp5040101

Frimpong, V. (2026). *AI debris: Residual risk and the afterlife of failed AI systems.* Social Science Research Network. https://ssrn.com/abstract=6313078

Gbabo, E. Y., Okenwa, O. K., & Chima, P. E. (2022). Framework for integrating cybersecurity risk controls into energy system implementation lifecycles. *Journal of Frontiers in Multidisciplinary Research, 3*(1), 365–371. https://doi.org/10.54660/jfmr.2022.001

Gupta, P. (2025). AI TIPS 2.0: A comprehensive framework for operationalizing AI governance. *arXiv.* https://doi.org/10.48550/arxiv.2512.09114

Hermansyah, H., & Zakaria, R. M. (2024). Integrating governance, technical controls, and agile practices: A multi-layered risk management framework for high technology projects. *Novatio: Journal of Management Technology and Innovation, 2*(3), 189–204. https://doi.org/10.61978/novatio.v2i3.851

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75–105. https://doi.org/10.2307/25148625

Hong, K. S., Chi, Y. P., Chao, L. R., & Tang, J. H. (2003). An integrated system theory of information security management. *Information Management & Computer Security, 11*(5), 243–248. https://doi.org/10.1108/09685220310500153

Hundeyin, W. O., Adegbenro, S. A., Rindap, Y. P., & Adaba, C. A. (2025). Integrating privacy-preserving AI models into AI governance frameworks. *International Journal of Innovative Science and Research Technology, 10*(10), 376–384. https://doi.org/10.38124/ijisrt/25oct209

Jooda, T. O., & Onukak, P. I. (2023). AI-driven governance, risk and compliance (GRC) systems: Enhancing cybersecurity resilience in financial institutions. *Open Access Research Journal of Science and Technology, 9*(1), 087–101. https://doi.org/10.53022/oarjst.2023.9.1.0054

Kabir, M. H., Razib, M., Arafat, Y., Rashed, R. A. M., & Jesan, Z. (2025). Strengthening U.S. critical infrastructure resilience through NIST-aligned cybersecurity governance and AI-driven threat detection. *Journal of Computer Science and Technology Studies, 7*(1), 1–15. https://doi.org/10.32996/jcsts.2025.7.1.1

Katakam, R. (2025). A unified model for ethical, transparent, and explainable AI in large-scale organizations. *Journal of Information Systems Engineering and Management, 10*(63s), 420–432. https://doi.org/10.52783/jisem.v10i63s.13879

Kemp, R. (2025). Security audits on artificial intelligence systems. *Cyber Security: A Peer-Reviewed Journal, 9*(1), 72. https://doi.org/10.69554/dfce5797

Khalid, H., & Jaffery, H. (2025). Managing the product lifecycle for AI-enabled products: From prototype to responsible retirement. *International Multidisciplinary Journal of Science, Technology & Business, 4*(3), Article 202329. https://imjstb.com/index.php/Journal/article/view/100

Lawrence, C., Cui, I., & Ho, D. E. (2023). The bureaucratic challenge to AI governance: An empirical assessment of implementation at U.S. federal agencies. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. https://doi.org/10.1145/3600211.3604701

Leon, M. (2026). Lifecycle-based governance to build reliable ethical AI systems. *Systems Research and Behavioral Science.* Advance online publication. https://doi.org/10.1002/sres.70014

Lestari, M., Wijaya, A. F., Sari, M. K., & Leander, L. K. (2025). Artificial intelligence governance for innovation and resilient data security in academic institutions. *Proceedings of the 2025 International Conference on ICT for Management in Complex Systems,* 1238–1243. https://doi.org/10.1109/icimcis68501.2025.11326907

Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry.* Sage. https://doi.org/10.4135/9781412986335

Loucif, S., Sharma, R., Kshetri, N., & Zahid, A. (2025). From design to decommissioning: TAFES framework for responsible AI. *Law, Ethics & Technology. Advance online publication*. https://doi.org/10.55092/let20250009

March, S. T., & Smith, G. F. (1995). Design and natural science research on information technology. *Decision Support Systems, 15*(4), 251–266. https://doi.org/10.1016/0167-9236(94)00041-2

Melaku, H. M. (2023). A dynamic and adaptive cybersecurity governance framework. *Journal of Cybersecurity and Privacy, 3*(4). https://doi.org/10.3390/jcp3040027

Motiang, N. C., & Mamorobela, S. P. (2025). Knowledge sharing framework for enhancing information security in South Africa's banking sector. In *5th International Multidisciplinary Information Technology and Engineering Conference (IMITEC). IEEE.* https://doi.org/10.1109/IMITEC67386.2025.11410455

Mwanje, D., Samuel, O., Tumwebaze, G., & Bukenya, M. (2023). A framework to enhance information security governance in SMEs. *Saudi Journal of Engineering and Technology, 8*(12), 300–303. https://doi.org/10.36348/sjet.2023.v08i12.002

National Institute of Standards and Technology. (2020). *Security and privacy controls for information systems and organizations (NIST Special Publication 800-53, Rev. 5).* U.S. Department of Commerce. https://doi.org/10.6028/NIST.SP.800-53r5

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1).* U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1

Nicho, M. (2018). A process model for implementing information systems security governance. *Information and Computer Security, 26*(1), 10–38. https://doi.org/10.1108/ICS-07-2016-0061

Nwinyi, I. P., & Oman-Amoako, M. (2026). AI-driven business analytics frameworks for government financial management: Advancing accountability and transparency in U.S. federal programs. *EPRA International Journal of Economics, Business and Management Studies.* Advance online publication. https://doi.org/10.36713/epra26112

Obimakinde, K. (2026). *The AI sunset protocol: A socio-technical framework for high-dependency AI decommissioning.* Social Science Research Network. https://ssrn.com/abstract=6157026

Odedina, E. (2024). AI-powered GRC: Enhancing regulatory compliance and risk resilience in evolving cyber threat landscapes. *World Journal of Advanced Research and Reviews, 23*(2), 3281–3290. https://doi.org/10.30574/wjarr.2024.23.3.2778

Owolabi, B. (2026). *Ethical and governance challenges of AI-based decision intelligence systems*. Social Science Research Network. https://doi.org/10.2139/ssrn.5085850

Paul Greene, F. (2025). Govern once/comply many: Leveraging cyber security framework experience to support AI governance. *Cyber Security: A Peer-Reviewed Journal, 9*(1), 49–60. https://doi.org/10.69554/vblf1901

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems, 24*(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Rahayu, I. S., Mulyana, R., & Fakhrurroja, H. (2025). Ambidextrous IoT governance framework for SmartCo's digital transformation aligned with COBIT 2019 traditional and DevOps. *Jurnal Nasional Pendidikan Teknik Informatika: JANAPATI, 14*(2), 308–320. https://doi.org/10.23887/janapati.v14i2.98135

Rahman, M. F., Ehsan, S. B., Rahman, T., & Mostafa, M. (2026). A FedRAMP and NIST aligned framework for securing AI systems in government clouds. *International Journal of Innovative Science and Research Technology, 11*(2), 911–918. https://doi.org/10.38124/ijisrt/26feb600

Ramli, A., & Darus, M. Y. (2026). AI-ready ICT security for education: A holistic framework to strengthen data integrity at Malaysia's data centre. *International Journal on E-Learning and Higher Education, 21*(1), 121–139. https://doi.org/10.24191/ijelhe.v21n1.2117

Rashid, S., Gurung, D., Gupta, S., & Rath, S. (2026). Lifecycle-integrated security for AI-cloud convergence in cyber-physical infrastructure. *arXiv*. https://doi.org/10.48550/arxiv.2602.23397

Rathnasinghe, A. P., Rahubadda, A. D., Ede, K. A., & Gledson, B. J. (2026). Codify, condition, capacitate: Expert perspectives on institution-first blockchain–BIM governance for PPP transparency in Nigeria. *FinTech, 5*(1), 10. https://doi.org/10.3390/fintech5010010

Reuel, A., & Undheim, T. (2024). Generative AI needs adaptive governance. *arXiv*. https://arxiv.org/abs/2405.03936

Robles, P., & Mallinson, D. J. (2026). Regulating AI in state governments: Security challenges and legislative responses. *International Journal of Public Administration, 49*(2), 71–84. https://doi.org/10.1080/01900692.2025.2598029

Saleh, K., & Abdulsalam, H. (2025). Operationalizing ISO/IEC 42001: Requirements and conformance evidence for AI management systems. In *2025 International Conference on Artificial Intelligence for Sustainable Innovation (AI-SI). IEEE.* https://doi.org/10.1109/AI-SI66213.2025.11341700

Saurabh, K. (2026). Strategic program management models for secure AI adoption in critical IT infrastructure. *The American Journal of Engineering and Technology, 8*(1), 87–96. https://doi.org/10.37547/tajet/volume08issue01-12

Sudarsan, S., Mittal, A., & Chandrasekaran, A. S. (2025). Secure AI-SDLC for critical infrastructure: Operationalizing the NIST AI RMF with evidence-driven controls. In *2025 International Conference on Computer and Applications (ICCA). IEEE.* https://doi.org/10.1109/ICCA66035.2025.11430939

Tremblay, M. C., Hevner, A. R., & Berndt, D. J. (2010). Focus groups for artifact refinement and evaluation in design research. *Communications of the Association for Information Systems, 26,* Article 27. https://doi.org/10.17705/1CAIS.02627

Ullah, K. I. (2024). Operationalizing information security governance: From framework adoption to control effectiveness. *Research Corridor Journal of Engineering Science, 1*(2). https://doi.org/10.66320/jrs99q91

Urhobo, B. (2024). Understanding the role of artificial intelligence in enhancing GRC practices in cybersecurity. *World Journal of Advanced Research and Reviews, 22*(2), 269–274. https://doi.org/10.30574/wjarr.2024.22.2.1340

Veiga, A. D., & Eloff, J. H. P. (2007). An information security governance framework. *Information Systems Management, 24*(4), 361–372. https://doi.org/10.1080/10580530701586136

Venable, J., Pries-Heje, J., & Baskerville, R. (2012). A comprehensive framework for evaluation in design science research. In *Proceedings of the 7th International Conference on Design Science Research in Information Systems and Technology* (pp. 423–438). Springer. https://doi.org/10.1007/978-3-642-29863-9_31

Venable, J., Pries-Heje, J., & Baskerville, R. (2016). FEDS: A framework for evaluation in design science research. *European Journal of Information Systems, 25*(1), 77–109. https://doi.org/10.1057/ejis.2014.21

Wijaya, A. F., Lestari, M., Sari, M. K., & Ferdiandinata, R. (2026). Public sector innovation in smart cities: A framework of ambidextrous AI governance. *JITK (Jurnal Ilmu Pengetahuan dan Teknologi Komputer), 11*(3), 765–774. https://doi.org/10.33480/jitk.v11i3.7672

Willie, M. (2025). Artificial intelligence and automation in administrative practices: A critical examination of the impact on institutional memory and governance frameworks. *Advances in Human Resource Development and Management, 1*(1), 25–35. https://ojs.nexuspress.org/journal-ahrdm/article/view/37

Yaokumah, W. (2017). Modelling the impact of administrative access controls on technical access control measures. *Information Resources Management Journal, 30*(4), 53–70. https://doi.org/10.4018/IRMJ.2017100104

Yasin, M., Arman, A. A., Edward, I. J. M., & Shalannanda, W. (2020). Designing information security governance recommendations and roadmap using COBIT 2019 framework and ISO 27001:2013. In 2020 *14th International Conference on Telecommunication Systems, Services, and Applications (TSSA). IEEE.* https://doi.org/10.1109/TSSA51342.2020.9310875

Appendices

Appendix A

Lifecycle-Integrated AI Governance Control Framework (LIAGCF)

Table A1 Control-Based AI Lifecycle Mapping of Governance Control Categories to NIST AI Risk Management Framework Functions and NIST SP 800-53 Control Families

Control Category	Control Objective	NIST AI RMF Function	NIST SP 800-53 Control Family	AI Lifecycle Phase	Example Governance Artifacts
Administrative Controls — GOVERN and MAP Functions					
Administrative	Establish a formal enterprise AI governance policy and scope	GOVERN	PM (Program Management), PL (Planning)	Strategy & Design	AI Governance Charter; Enterprise AI Policy
Administrative	Define AI accountability, decision rights, and risk ownership	GOVERN	PM (Program Management)	Strategy & Design	AI RACI Matrix; Governance Role Definitions
Administrative	Establish third-party AI risk governance policies	GOVERN	SR (Supply Chain Risk Management), SA (System & Services Acquisition)	Procurement	Vendor AI Risk Review Checklist; Third-Party Risk Policy
Administrative	Define enterprise AI audit and governance review cadence	GOVERN	CA (Assessment, Authorization & Monitoring), PM (Program Management)	Operations	AI Governance Review Plan; Audit Schedule
Administrative	Require AI risk assessments prior to deployment	MAP	RA (Risk Assessment), CA (Assessment, Authorization & Monitoring)	Development	AI Risk Assessment Report; Risk Register Entry
Administrative	Categorize AI systems by risk level, use context, and affected population	MAP	RA (Risk Assessment), PL (Planning)	Strategy & Design	AI System Inventory; Risk Categorization Register
Technical Controls — MAP and MEASURE Functions					
Technical	Identify and assess potential for discriminatory or disparate impact outcomes	MAP	RA (Risk Assessment), SI (System & Information Integrity)	Development	Bias Risk Assessment; Fairness Evaluation Report
Technical	Map data provenance, lineage, and quality risks for training datasets	MAP	SA (System & Services Acquisition), SR (Supply Chain Risk Management)	Data Preparation	Data Lineage Map; Training Data Quality Report
Technical	Validate AI models prior to production release	MEASURE	CA (Assessment, Authorization & Monitoring), SA (System & Services Acquisition)	Development	Model Validation Report; Testing Summary
Technical	Monitor model performance and detect drift	MEASURE	SI (System & Information Integrity), CA (Assessment, Authorization & Monitoring)	Operations	Drift Monitoring Dashboard; Performance Metrics Report
Technical	Log AI system activity for traceability and auditability	MEASURE	AU (Audit & Accountability)	Operations	Audit Logs; Monitoring Dashboard
Operational Controls — MEASURE and MANAGE Functions					
Operational	Conduct periodic AI governance and risk reviews	MEASURE	CA (Assessment, Authorization & Monitoring), RA (Risk Assessment)	Operations	Governance Review Report; Risk Reassessment Summary
Operational	Enforce access controls on AI training and operational data	MANAGE	AC (Access Control), IA (Identification & Authentication)	Data Preparation / Operations	Access Control Matrix; Privileged Access Review
Operational	Protect AI models from unauthorized modification	MANAGE	CM (Configuration Management), SI (System & Information Integrity)	Operations	Configuration Baseline; Change Control Log
Operational	Execute AI incident response playbooks	MANAGE	IR (Incident Response), SI (System & Information Integrity)	Operations	AI Incident Playbook; After-Action Report
Operational	Retire AI systems exceeding defined risk tolerance	MANAGE	SA (System & Services Acquisition), CM (Configuration Management), MP (Media Protection)	Retirement	AI Retirement Plan; Risk Closure Documentation

Note. Control objectives are illustrative and designed to support governance development and assurance alignment. Organizations should tailor control selection and implementation based on risk appetite, regulatory environment, and system importance. Color coding (where rendered): blue = Administrative controls; green = Technical controls; yellow = Operational controls. Row grouping headers provide an alternative visual differentiator if color is not preserved in production.

GOVERN = Organizational accountability and policy; MAP = Risk identification and categorization; MEASURE = Assessment and evaluation; MANAGE = Operational enforcement and lifecycle management.

Appendix B

Expert Assessment Survey Instrument and RQ-to-Survey-Question Mapping

Study: Evaluating Control-Based AI Governance in Cybersecurity GRC Programs: An Expert Assessment Study

Researcher: Miranda Stanfield | stanfield.miranda@gmail.com | Capitol Technology University, PhD Cybersecurity Leadership

Survey Platform: Alchemer | Survey URL: https://survey.alchemer.com/s3/8712601/

A. Part 1: RQ-to-Survey-Question Mapping

Table B1 shows the explicit mapping between the five research questions and the five survey questions given to experts. Each question was designed to gather specific judgments needed to evaluate the corresponding research question. This mapping aligns the evaluation instrument with the study's objectives, allowing findings in Section 5 to be directly linked to the research questions in Section 1.3.

Table B1 RQ-to-Survey-Question Mapping: Expert Assessment of the LIAGCF

RQ	Survey Q	Research Question	Survey Question Text (Abbreviated)	Evaluative Purpose
RQ1	**Q6**	How effectively do the LIAGCF governance controls align with the four NIST AI RMF core functions, GOVERN, MAP, MEASURE, and MANAGE?	In your expert judgment, how effectively does the LIAGCF align governance controls across the four NIST AI RMF core functions — GOVERN, MAP, MEASURE, and MANAGE — and where, if anywhere, do you identify misalignments between the stated control objectives and their assigned RMF functions?	Elicit expert judgment on RMF function alignment and identify specific misalignment concerns
RQ2	**Q7**	How appropriately do the NIST SP 800-53 control family assignments operationalize each governance objective within the LIAGCF?	Examining the NIST SP 800-53 control family assignments within the framework, how appropriately do the assigned control families operationalize each governance objective, and are there specific rows where you believe the control family mapping should be revised, expanded, or replaced to better reflect AI governance requirements?	Assess SP 800-53 operationalization quality and surface specific revision recommendations
RQ3	**Q8**	How comprehensively does the LIAGCF represent governance requirements across the seven AI system lifecycle phases from Strategy and Design through Retirement, and what governance domains or risk considerations remain insufficiently addressed?	Considering the full AI system lifecycle from Strategy & Design through Retirement, how comprehensively does the framework represent governance controls across all critical lifecycle phases, and are there phases, risk domains, or governance concerns — such as fairness, transparency, or sociotechnical harm — that you believe are insufficiently addressed?	Evaluate lifecycle coverage completeness and identify underrepresented phases or risk domains
RQ4	**Q9**	Do expert practitioners consider the LIAGCF's integration of NIST AI RMF and SP 800-53 within a lifecycle-aligned structure to represent a meaningful and distinctive contribution to AI governance practice and scholarship?	In your assessment, does the integration of the NIST AI Risk Management Framework and NIST SP 800-53 within a unified, lifecycle-aligned governance control structure represent a meaningful and novel contribution to the AI governance literature, and what do you consider to be the framework's most significant strength and the most critical revision needed before scholarly publication or operational deployment?	Establish expert-validated face and construct validity; identify most significant strength and critical revision needed
RQ5	**Q10**	What operational barriers would organizations encounter when integrating the LIAGCF into existing GRC infrastructure?	From a practitioner perspective, how effectively could this framework be integrated into an organization's existing governance, risk, and compliance infrastructure — including established enterprise risk management programs, cybersecurity frameworks, and AI oversight mechanisms? What do you identify as the most significant operational barriers an organization would likely encounter when attempting to adopt and sustain this framework in practice?	Assess practical integration feasibility and identify the most significant adoption barriers

Note. RQ = Research Question. Q = Survey Question. All five substantive questions are open-ended. Abbreviations: LIAGCF = Lifecycle-Integrated AI Governance Control Framework; AI RMF = NIST Artificial Intelligence Risk Management Framework; SP 800-53 = NIST Special Publication 800-53, Security and Privacy Controls for Information Systems and Organizations; GRC = Governance, Risk, and Compliance. The Q9 instrument text uses "novel contribution to the AI governance literature" as the operational phrasing of the RQ4 construct "distinctive contribution to AI governance practice and scholarship"; "novel" was selected to elicit contribution-originality judgments from practitioners, while "distinctive" is the analytical term used in the research question.

B. Part 2: Eligibility Screening and Confirmation Gate Questions

Table B2 describes the five eligibility screening questions and the confirmation gate question (Q1–Q5) that precede the substantive evaluation instrument. These questions serve three functions: (1) confirming participant eligibility for inclusion in the study sample; (2) capturing demographic and expertise data used in the expert panel profile reported in Section 5.1; and (3) ensuring that all substantive responses are based on direct review of the LIAGCF artifact.

Table B2 Eligibility Screening and Confirmation Gate Questions: Function and Data Use

Q#	Question Label	Response Type	Inclusion/Exclusion Function	Data Use in Study
Q1	Consent & Eligibility	Checkbox (4 conditions)	Confirms all eligibility criteria are met; participant excluded if not all checked	Used to validate sample integrity; not reported in demographic tables
Q2	Years of Experience	Single-select (4 options)	Captures experience level; minimum 5 years required for inclusion	Reported in Section 5.1 panel profile and Appendix C, Table C1
Q3	Primary Expertise	Multi-select + write-in	Documents domain expertise profile of expert panel; used in Section 5.1 demographics	Reported in Section 5.1 panel profile and Appendix C, Table C2
Q4	NIST Framework Familiarity	Single-select per framework (SP 800-53 and AI RMF)	Captures framework familiarity; AI RMF familiarity distribution reported as limitation	Reported in Section 5.1 panel profile and Appendix C, Table C3; AI RMF distribution discussed as limitation in Section 6.7
Q5	Table Review Confirmation	Single-select (required gate)	Ensures all substantive responses are based on direct artifact review; non-confirmers excluded	Used to validate that analysis dataset contains only artifact-informed responses; not reported in demographic tables

Note. Q = Survey Question number as administered in Alchemer. Responses to Q1–Q4 were used to construct the expert panel demographic profile reported in Section 5.1. Q5 is a required gate question; responses submitted without confirming table review were excluded from the analysis dataset.

C. Part 3: Complete Survey Instrument

The following presents the complete survey instrument as administered to participants via Alchemer. Question numbering for Q1–Q10 corresponds to the mappings in Tables B1 and B2 above. Questions 11 and 12 are optional and administrative and are not mapped to research questions. All question text is reproduced exactly as presented to respondents.

D. Section 1: Eligibility Screening

The following questions confirm that you meet the eligibility criteria for participation in this study. All four conditions in Question 1 must be satisfied to proceed.

Question 1 — Consent and Eligibility Confirmation (Required)

Please confirm that all of the following apply to you:

- ☐ You are at least 18 years old
- ☐ You have at least 5 years of professional experience in cybersecurity, risk management, compliance, or AI governance
- ☐ You have working familiarity with NIST-based governance frameworks
- ☐ You voluntarily agree to participate

Response options: I have read the information above and agree to participate. | I do not agree to participate.

Question 2 — Years of Professional Experience

How many years of professional experience do you have in cybersecurity, risk, or compliance?

Response options: 5–10 years | 11–15 years | 16–20 years | 20+ years

Question 3 — Primary Area(s) of Expertise (Select all that apply)

- ☐ Cybersecurity Governance
- ☐ Enterprise Risk Management
- ☐ Compliance & Audit
- ☐ Federal GRC Programs
- ☐ AI/ML Security
- ☐ Commercial GRC Program
- ☐ Other (please specify): _____

Question 4 — NIST Framework Familiarity

Please indicate your current level of familiarity with each framework:

NIST SP 800-53:

Extensive experience (have implemented or assessed controls) | Working familiarity (can read and apply) | Limited familiarity (general awareness) | No familiarity

NIST AI RMF (AI 100-1):

Extensive experience (have implemented or assessed) | Working familiarity (can read and apply) | Limited familiarity (general awareness) | No familiarity

*1) **Section 2: Table Review Confirmation***

Before proceeding to the substantive evaluation questions, you must confirm that you have reviewed the LIAGCF governance control mapping table in full. Responses from participants who have not reviewed the table will be excluded from analysis.

Question 5 — Confirmation Gate (Required)

Please confirm the following before proceeding:

Response options: I have reviewed and understand how to read and interpret the governance control mapping table above. | I have not yet reviewed the table and cannot proceed at this time.

*2) **Section 3: Expert Evaluation of the LIAGCF (Questions 6–10)***

The following five questions constitute the substantive expert evaluation of the LIAGCF. Each question corresponds directly to one of the five research questions guiding this study, as mapped in Table B1 above. All responses are open-ended. Please respond based on your direct review of the governance control mapping table provided.

Question 6 — Framework Alignment

Maps to RQ1 | Evaluates structural alignment of 16 controls with NIST AI RMF core functions

In your expert judgment, how effectively does the Lifecycle-Integrated AI Governance Control Framework (LIAGCF) align governance controls across the four NIST AI RMF core functions — GOVERN, MAP, MEASURE, and MANAGE — and where, if anywhere, do you identify misalignments between the stated control objectives and their assigned RMF functions?

[Open-ended response field]

Question 7 — SP 800-53 Operationalization

Maps to RQ2 | Evaluates appropriateness of SP 800-53 control family assignments

Examining the NIST SP 800-53 control family assignments within the framework, how appropriately do the assigned control families operationalize each governance objective, and are there specific rows where you believe the control family mapping should be revised, expanded, or replaced to better reflect AI governance requirements?

[Open-ended response field]

Question 8 — AI Lifecycle Representation

Maps to RQ3 | Evaluates comprehensiveness of lifecycle phase coverage

Considering the full AI system lifecycle from Strategy & Design through Retirement, how comprehensively does the framework represent governance controls across all critical lifecycle phases, and are there phases, risk domains, or governance concerns — such as fairness, transparency, or sociotechnical harm — that you believe are insufficiently addressed?

[Open-ended response field]

Question 9 — Theoretical Contribution

Maps to RQ4 | Evaluates perceived novelty and scholarly contribution of the LIAGCF

In your assessment, does the integration of the NIST AI Risk Management Framework and NIST SP 800-53 within a unified, lifecycle-aligned governance control structure represent a meaningful and novel contribution to the AI governance literature, and what do you consider to be the framework's most significant strength and the most critical revision needed before scholarly publication or operational deployment?

[Open-ended response field]

Question 10 — Operational Integration

Maps to RQ5 | Evaluates practical integration feasibility and identifies adoption barriers

From a practitioner perspective, how effectively could this framework be integrated into an organization's existing governance, risk, and compliance infrastructure — including established enterprise risk management programs, cybersecurity frameworks, and AI oversight mechanisms? What do you identify as the most significant operational barriers an organization would likely encounter when attempting to adopt and sustain this framework in practice?

[Open-ended response field]

3) Section 4: Optional Comments and Follow-Up

Question 11 — Additional Comments (Optional)

Is there anything about the framework's structure, scope, methodology, or presentation that you would like to comment on that has not been addressed by the preceding questions?

[Open-ended response field — optional]

Question 12 — Willingness to Clarify

Would you be willing to clarify your responses if the researcher has follow-up questions?

Response options: Yes (please provide email) | No

[Email address — optional]

End of Survey Instrument. Thank you for your participation.
All responses are confidential and will be reported only in anonymized, aggregate form.

Appendix C

Expert Panel Demographic Summary (N = 10)

Study: Evaluating Control-Based AI Governance in Cybersecurity GRC Programs: An Expert Assessment Study

This appendix shows the demographic profile of the 10 expert respondents whose responses form the dataset. All data are in aggregate; no individual info is shown. Three PDF exports had blank Q5 fields due to Alchemer platform issues on mobile submissions; those responses are excluded. Total N = 10 includes confirmed complete responses meeting all criteria. Table C1 reports the experience distribution, Table C2 reports the expertise area distribution, Table C3 reports NIST framework familiarity levels, and Table C4 reports follow-up clarification consent.

Table C1 Expert Panel Years of Professional Experience Distribution

Years of Professional Experience	n	%
5–10 years	5	50%
11–15 years	3	30%
16–20 years	2	20%
Total	10	100%

Note. Experience levels reflect self-reported responses to Q2 of the survey instrument. All respondents meet the minimum eligibility criterion of five or more years of professional experience in cybersecurity, risk management, compliance, or AI governance.

Table C2 Expert Panel Primary Expertise Area Distribution

Primary Expertise Area (select all that apply)	n	% of Panel
Cybersecurity Governance	8	80%
Compliance & Audit	7	70%
Federal GRC Programs	7	70%
Enterprise Risk Management	6	60%
Commercial GRC Program	4	40%
AI/ML Security	3	30%

Note. Expertise areas reflect multiple-selection responses to Q3 of the survey instrument. Percentages reflect the proportion of the N = 10 panel selecting each area and sum to more than 100%.

Table C3 Expert Panel NIST Framework Familiarity Distribution

Familiarity Level	Description	SP 800-53 n (%)	AI RMF n (%)
Extensive experience	Has implemented or assessed controls	5 (50%)	1 (10%)
Working familiarity	Can read and apply the framework	5 (50%)	6 (60%)
Limited familiarity	General awareness only	0 (0%)	2 (20%)
No familiarity	No prior exposure	0 (0%)	1 (10%)
Total		10 (100%)	10 (100%)

Note. Familiarity levels reflect self-reported responses to Q4 of the survey instrument. SP 800-53 = NIST Special Publication 800-53. AI RMF = NIST Artificial Intelligence Risk Management Framework. The AI RMF familiarity distribution, with three respondents reporting limited or no familiarity, is acknowledged as a study limitation in Section 6.7.

Table C4 Expert Panel Follow-Up Clarification Consent

Follow-Up Willingness	n	%
Willing to clarify (provided email)	4	40%
Willing to clarify (no email provided)	2	20%
Not willing or no response	4	40%

Note. Follow-up willingness reflects responses to Q12 of the survey instrument. No follow-up contact was initiated as part of this study. Email addresses collected are not reported.

The expert panel's demographic profile aligns with the LIAGCF's target audience of federal GRC practitioners, mainly in Cybersecurity Governance (80%), Compliance and Audit (70%), and Federal GRC Programs (70%), covering key organizational contexts for deployment. The split between extensive and working SP 800-53 familiarity shows the panel includes both deep implementers and applied practitioners, the main LIAGCF users. The varied AI RMF familiarity reflects its current adoption level and highlights the practitioner gap the LIAGCF is designed to address.

Appendix D

Thematic Analysis Codebook Summary

Study: Evaluating Control-Based AI Governance in Cybersecurity GRC Programs: An Expert Assessment Study

This appendix presents a summary codebook from a six-phase thematic analysis by Braun and Clarke (2006) applied to expert responses across evaluation dimensions (Q6–Q10). The codebook documents themes from Phases 3–5: initial code grouping, candidate review, and theme definition. It records theme code, name, definition, initial codes, and prevalence among N=10 experts.

The researcher maintains the full analytic audit trail, including raw codes mapped to response segments, which is available upon request. This summary codebook adheres to transparency standards for qualitative DSR research by Rathnasinghe et al. (2026) and Lincoln and Guba (1985).

Table D1 Thematic Analysis Codebook Summary: Final Theme Set Organized by Evaluation Dimension (N = 10)

Theme Code	Theme Name	Theme Definition	Representative Initial Codes	Prevalence
RQ1 — Framework Alignment with NIST AI RMF Core Functions (Q6)				
T1-A	**Structural Alignment Broadly Validated**	The LIAGCF's distribution of controls across GOVERN, MAP, MEASURE, and MANAGE was confirmed as logically sequenced and consistent with NIST AI RMF 1.0 functional intent, with GOVERN appropriately anchoring policy and accountability.	ALIGNMENT_VALIDATED; GOVERN_FOUNDATION_VALIDATED; TOP_DOWN_LOGIC_VALIDATED	9 of 10 — DOMINANT
T1-B	**Cross-Functional Control Behavior at Function Boundaries**	Certain controls — particularly bias assessment and data provenance — exhibit functional characteristics of both MAP and MEASURE. Respondents characterized this as inherent to integrated governance frameworks, not a structural flaw.	BOUNDARY_BLUR_MAP_MEASURE; INHERENT_OVERLAP_EXPECTED	6 of 10 — STRONG
T1-C	**MANAGE/MEASURE Boundary Clarification Needed**	A minority of respondents noted that audit logging controls assigned to MEASURE carry GOVERN-level accountability characteristics, and that MANAGE risks becoming a catch-all category.	MANAGE_OPERATIONALIZATION_GAP; AUDIT_LOGGING_PLACEMENT	2 of 10 — MODERATE
RQ2 — SP 800-53 Control Family Operationalization (Q7)				
T2-A	**SP 800-53 Operationalization Broadly Validated**	Expert respondents confirmed that SP 800-53 control family assignments provide a defensible operational foundation. PM, PL, RA, CA, AU, IR, and CM were consistently identified as well-matched.	OPERATIONALIZATION_VALIDATED; STRONG_FAMILIES_CONFIRMED	7 of 10 — DOMINANT

Theme Code	Theme Name	Theme Definition	Representative Initial Codes	Prevalence
T2-B	**Targeted Expansion Recommendations for AI-Specific Risk**	Respondents identified specific rows warranting expansion: PT for the bias/fairness control; AU for the Retirement control; CA for formal fairness evaluation validation; SI for training data integrity.	PT_FAMILY_RECOMMENDED; AU_RETIREMENT_RECOMMENDED; BIAS_FAIRNESS_EXPANSION; TRAINING_DATA_EXPANSION; THIRD_PARTY_EXPANSION	7 of 10 — DOMINANT
T2-C	**SP 800-53 AI-Native Tension**	One respondent noted SP 800-53's design for IT security creates tension with AI-native governance requirements (model access governance, foundation model provenance), recommending hybrid supplementation with NIST AI 600-1.	SP80053_AI_TENSION	1 of 10 — NOTABLE
RQ3 — AI Lifecycle Representation (Q8)				
T3-A	**Seven-Phase Coverage Comprehensive**	Expert respondents validated the seven-phase lifecycle structure as comprehensive. Data Preparation, Procurement, and Retirement phases received specific recognition as underrepresented in comparable frameworks.	LIFECYCLE_COMPREHENSIVE; DATA_PREPARATION_STRENGTH; PROCUREMENT_STRENGTH; RETIREMENT_STRENGTH	7 of 10 — DOMINANT
T3-B	**Transparency, Explainability, and Post-Deployment Fairness as Priority Gaps**	Five respondents identified the absence of explicit controls for model transparency, explainability documentation, and ongoing fairness monitoring in Operations as the most significant lifecycle gap.	TRANSPARENCY_EXPLAINABILITY_GAP; FAIRNESS_MONITORING_OPS_GAP	5 of 10 — STRONG
T3-C	**Deployment Phase Governance Underrepresented**	The development-to-production transition lacks explicit controls for deployment authorization, production readiness review, rollback criteria, and human oversight thresholds.	DEPLOYMENT_GAP	2 of 10 — MODERATE
T3-D	**Retraining Governance and Human Oversight Undertheorized**	Absence of governance structures for retraining approval, dataset refresh governance, model re-certification, and human accountability when AI behavior falls outside anticipated parameters.	RETRAINING_GOVERNANCE_GAP; HUMAN_OVERSIGHT_GAP	2 of 10 — MODERATE
RQ4 — Theoretical Contribution (Q9)				
T4-A	**Consensus on Meaningful Scholarly and Practical Contribution**	Nine of ten respondents affirmed the LIAGCF as a meaningful contribution. Multiple respondents identified the framework as ready for operational deployment. The contribution resolves the gap between principled AI risk management and operational cybersecurity enforcement.	CONTRIBUTION_VALIDATED; OPERATIONAL_NOW	9 of 10 — DOMINANT
T4-B	**Audit Traceability and Translational Value as Primary Strength**	Eight respondents identified the framework's ability to translate AI governance intent into auditable, enforceable controls traceable from governance objective through lifecycle phase to implementation artifact as its primary and distinctive strength.	AUDIT_TRACEABILITY_STRENGTH; BRIDGES_TWO_TRACKS; ARTIFACTS_AS_STRENGTH	8 of 10 — DOMINANT
T4-C	**Contribution Is Integrative-Structural, Not Conceptually Originary**	Two respondents explicitly characterized the LIAGCF's contribution as synthesis and operationalization of existing NIST standards — a legitimate and valuable scholarly contribution that should be framed as such to withstand peer review scrutiny.	INTEGRATIVE_NOT_ORIGINARY	2 of 10 — MODERATE
T4-D	**Control Selection Methodology Documentation as Critical Pre-Publication Revision**	Four respondents identified the absence of explicit documentation for how the 16 controls were selected and how SP 800-53 families were assigned as the single most critical pre-publication revision.	METHODOLOGY_TRANSPARENCY_NEEDED	4 of 10 — STRONG
RQ5 — Operational Integration (Q10)				
T5-A	**Integration Feasible for NIST-Aligned Organizations**	Eight respondents confirmed integration feasibility, with SP 800-53 grounding identified as the primary enabler providing a natural on-ramp for organizations already running NIST-aligned programs.	INTEGRATION_FEASIBLE; NIST_ALIGNMENT_ENABLES_INTEGRATION	8 of 10 — DOMINANT
T5-B	**Organizational Barriers Dominate Over Technical Barriers**	Seven respondents identified three primary organizational barriers: unclear AI risk ownership across legal, security, and data science functions; incomplete AI system inventories including shadow AI; and workforce expertise gaps combining AI and GRC knowledge.	OWNERSHIP_AMBIGUITY_BARRIER; AI_INVENTORY_GAP_BARRIER; EXPERTISE_GAP_BARRIER	7 of 10 — DOMINANT
T5-C	**Technical Infrastructure Gaps as Secondary Barriers**	Three respondents noted that MEASURE controls assume AI-specific monitoring infrastructure — drift detection, bias monitoring, AI performance metrics — not yet available in most GRC platforms.	TOOLING_LIMITATION_BARRIER; MONITORING_INFRASTRUCTURE_GAP	3 of 10 — MODERATE
T5-D	**Maturity-Tiered Implementation Guidance Needed**	Three respondents observed the framework assumes organizational AI governance maturity and recommended tiered implementation pathways to broaden applicability across governance development stages.	MATURITY_VARIANCE_BARRIER; PARALLEL_TRACK_RISK	3 of 10 — MODERATE

Note. Prevalence designations: DOMINANT = identified by 7 or more respondents; STRONG = identified by 4–6 respondents; MODERATE = identified by 2–3 respondents; NOTABLE = identified by 1 respondent but carries significant theoretical weight. Themes are organized by evaluation dimension corresponding to RQ1–RQ5. Representative initial codes are rendered in SMALL_CAPS format reflecting the coding vocabulary used in Phase 2 of the analysis. Theme definitions reflect Phase 5 finalization following review against the full coded dataset in Phase 4.

The codebook's 18 themes across five evaluation dimensions provide the evidentiary structure for the findings in Section 5 and the interpretive analysis in Section 6. Eight dominant themes emerged, supported by seven or more respondents: structural alignment validated (T1-A), SP 800-53 operationalization validated (T2-A),

targeted expansion recommendations (T2-B), lifecycle coverage comprehensive (T3-A), consensus on contribution (T4-A), integration feasibility (T5-A), audit traceability as the primary strength (T4-B), and organizational barriers as the dominant adoption challenge (T5-B). These dominant themes collectively confirm the LIAGCF's structural soundness and identify the specific development priorities: transparency controls, post-deployment fairness monitoring, and tiered implementation guidance, that will advance the framework's scholarly rigor and operational applicability in future iterations.

Appendix E

NIST SP 800-53 Control Family Selection: Design Rationale for the LIAGCF's 16 Governance Controls

Table E1. Design Rationale for NIST SP 800-53 Control Family Assignments Across the 16 LIAGCF Governance Controls

Control #	Category	Control Objective	AI RMF Function	SP 800-53 Control Family Assignment	Design Rationale for Control Family Selection
Administrative Controls — GOVERN Function (Controls A-1 through A-4)					
A-1	Administrative	Establish a formal enterprise AI governance policy and scope	**GOVERN**	PM (Program Management), PL (Planning)	PM establishes the organizational infrastructure — program structure, roles, and accountability — required to sustain AI governance as an enterprise function. PL provides the planning framework that translates governance intent into documented policies, acceptable use definitions, and scope boundaries. Together, PM and PL operationalize the GOVERN function's requirement to establish organizational accountability and governance authority before any AI system is acquired or deployed.
A-2	Administrative	Define AI accountability, decision rights, and risk ownership	**GOVERN**	PM (Program Management)	PM directly governs the organizational roles, responsibilities, and decision authorities required to assign AI risk ownership and define escalation pathways. The PM family includes controls for establishing senior leadership accountability, designating program officials, and creating oversight structures — each of which is necessary for governance of AI risk to be institutionally enforceable rather than advisory. No other SP 800-53 control family addresses organizational accountability structures at the program level.
A-3	Administrative	Establish third-party AI risk governance policies	**GOVERN**	SR (Supply Chain Risk Management), SA (System & Services Acquisition)	SR addresses risks introduced through external AI components, vendor relationships, and supply chain dependencies — the primary third-party risk domain for AI procurement. SR controls establish due diligence requirements, supplier assessment processes, and contractual governance obligations. SA complements SR by governing acquisition processes and vendor evaluation criteria, ensuring that third-party AI components meet organizational governance requirements before procurement. SP 800-53 Rev 5 introduced SR as a distinct control family specifically because supply chain risks had become a critical, underaddressed governance domain.
A-4	Administrative	Define enterprise AI audit and governance review cadence	**GOVERN**	CA (Assessment, Authorization & Monitoring), PM (Program Management)	CA governs the formal assessment and authorization processes through which AI systems are reviewed against established governance criteria at defined intervals. CA controls establish audit schedules, ongoing authorization requirements, and continuous monitoring obligations — the structural elements of a governance review cadence. PM provides the program-level oversight framework that ensures review activities are resourced, scheduled, and reported within the enterprise governance architecture. Together, CA and PM establish the institutionalized oversight rhythm that distinguishes a governance program from ad hoc compliance activity.
Administrative Controls — MAP Function (Controls A-5 through A-6)					
A-5	Administrative	Require AI risk assessments prior to deployment	**MAP**	RA (Risk Assessment), CA (Assessment, Authorization & Monitoring)	RA directly governs the risk assessment process, establishing the methods, criteria, and documentation requirements for identifying and evaluating risk before system deployment. RA controls require organizations to assess threats, vulnerabilities, and potential impacts — the core activities of the MAP function's risk identification requirement. CA provides the authorization gate that ensures risk assessment findings are formally reviewed and that deployment authorization is conditioned on acceptable risk posture. The combination of RA and CA creates a structured pre-deployment governance checkpoint aligned with the MAP function.
A-6	Administrative	Categorize AI systems by risk level, use context, and affected population	**MAP**	RA (Risk Assessment), PL (Planning)	RA governs risk categorization processes, requiring organizations to classify systems according to potential impact and risk profile. AI system categorization requires the same risk-based assessment framework that RA controls establish — identifying the sensitivity of data processed, the criticality of decisions supported, and the populations affected. PL provides the planning structure within which categorization decisions are documented and reflected in governance policies and acceptable use definitions. Categorization at the strategy phase, before system acquisition or development, is the MAP function's foundational risk contextualization activity.
Technical Controls — MAP Function (Controls T-1 through T-2)					
T-1	Technical	Identify and assess potential for discriminatory or disparate impact outcomes	**MAP**	RA (Risk Assessment), SI (System & Information Integrity)	RA governs the identification and evaluation of risks — including algorithmic bias and disparate impact — as a systematic assessment activity. Bias risk assessment is structurally a risk identification exercise: it requires analyzing potential impacts on affected populations, documenting findings, and informing risk treatment decisions. SI governs system integrity, including the integrity of model outputs and the mechanisms needed to detect and respond to integrity failures such as biased predictions. SI controls support the technical implementation of bias testing mechanisms that generate the evidence required for RA-based assessment.
T-2	Technical	Map data provenance, lineage, and quality risks for training datasets	**MAP**	SA (System & Services Acquisition), SR (Supply Chain Risk Management)	SA governs the acquisition of data and system components, establishing requirements for documentation, quality standards, and vendor obligations that apply directly to training dataset sourcing. Training data often originates from third-party sources, making SR directly applicable: SR controls govern the provenance verification, integrity assurance, and supply chain documentation requirements for external data components. Together, SA and SR address the full scope of training data risk at the Data Preparation phase, where provenance and quality failures introduce governance risk that cannot be remediated after model training is complete.
Technical Controls — MEASURE Function (Controls T-3 through T-5)					

Control #	Category	Control Objective	AI RMF Function	SP 800-53 Control Family Assignment	Design Rationale for Control Family Selection
T-3	Technical	Validate AI models prior to production release	**MEASURE**	CA (Assessment, Authorization & Monitoring), SA (System & Services Acquisition)	CA governs the formal assessment and authorization process, including testing and evaluation activities that determine whether a system meets established requirements before production release. Model validation is structurally an authorization activity: it generates the evidence required to make a deployment authorization decision. SA governs system development and acquisition requirements, establishing the validation and testing criteria that models must satisfy before organizational acceptance. The CA-SA combination mirrors the MEASURE function's requirement to analyze and assess AI risk through systematic evaluation before deployment.
T-4	Technical	Monitor model performance and detect drift	**MEASURE**	SI (System & Information Integrity), CA (Assessment, Authorization & Monitoring)	SI governs the mechanisms for detecting unauthorized changes, anomalies, and integrity failures in information systems — including the performance degradation and distributional drift that characterize AI model decay in production. SI controls require organizations to monitor system integrity continuously and respond to integrity events. CA provides the ongoing authorization and continuous monitoring framework within which drift detection findings are assessed against governance thresholds and trigger re-authorization or remediation requirements. Together, SI and CA address the MEASURE function's continuous assessment obligation for deployed AI systems.
T-5	Technical	Log AI system activity for traceability and auditability	**MEASURE**	AU (Audit & Accountability)	AU is the SP 800-53 control family specifically designed to govern audit logging, log integrity, and the accountability mechanisms that make system activity traceable and reviewable. AU controls require organizations to capture sufficient event data to reconstruct AI system decisions and actions, protect log integrity, and retain audit records for defined periods. For AI systems, AU-based logging directly supports the MEASURE function's requirement to generate evidence of system behavior that can be independently verified. No other SP 800-53 control family addresses auditability and decision traceability as its primary purpose.
Operational Controls — MEASURE Function (Control O-1)					
O-1	Operational	Conduct periodic AI governance and risk reviews	**MEASURE**	CA (Assessment, Authorization & Monitoring), RA (Risk Assessment)	CA governs the ongoing assessment, authorization, and monitoring processes that constitute periodic governance review — the structured activities through which organizations evaluate whether deployed AI systems continue to operate within accepted risk parameters. RA provides the risk assessment methodology that drives the substantive content of periodic reviews, requiring organizations to reassess risk profiles, update risk registers, and validate that risk treatment decisions remain appropriate. The CA-RA combination operationalizes the MEASURE function's requirement to continuously evaluate AI system risk through systematic, documented review cycles.
Operational Controls — MANAGE Function (Controls O-2 through O-5)					
O-2	Operational	Enforce access controls on AI training and operational data	**MANAGE**	AC (Access Control), IA (Identification & Authentication)	AC governs the authorization and enforcement of access rights to information and systems — the direct mechanism for restricting who can read, modify, or delete AI training datasets and operational model components. AC controls require organizations to define least-privilege access policies and enforce separation of duties, which are essential protections for AI data integrity. IA establishes the identity verification mechanisms that ensure access control decisions are applied to correctly identified principals. Together, AC and IA address the MANAGE function's requirement to operationally enforce governance policies through runtime access restrictions.
O-3	Operational	Protect AI models from unauthorized modification	**MANAGE**	CM (Configuration Management), SI (System & Information Integrity)	CM governs the formal change control processes that prevent unauthorized modifications to systems, configurations, and software components — including AI model artifacts, weights, and pipeline code. CM controls require organizations to maintain configuration baselines, document authorized changes, and review change requests against governance criteria before implementation. SI provides the integrity monitoring mechanisms that detect unauthorized modifications and alert governance functions. The CM-SI combination creates a two-layer protection architecture: CM prevents unauthorized changes through process controls, while SI detects them through technical monitoring.
O-4	Operational	Execute AI incident response playbooks	**MANAGE**	IR (Incident Response), SI (System & Information Integrity)	IR governs the organizational processes for detecting, responding to, and recovering from security and operational incidents — the control family most directly applicable to AI system failures, misuse events, and governance breaches. IR controls require organizations to maintain incident response plans, train response personnel, and execute documented playbooks when incidents occur. SI provides the monitoring and integrity checking mechanisms that generate the alerts and anomaly signals that trigger incident response. Together, IR and SI address the MANAGE function's requirement to actively respond to and manage AI risks as they materialize in production environments.
O-5	Operational	Retire AI systems exceeding defined risk tolerance	**MANAGE**	SA (System & Services Acquisition), CM (Configuration Management), MP (Media Protection)	SA governs system decommissioning and the formal processes through which systems are removed from authorized operation — the acquisition lifecycle control that covers retirement decisions and closure documentation. CM governs the configuration management activities required at retirement: removing system components from baselines, revoking change control records, and documenting final system states. MP governs the secure handling, sanitization, and disposal of media containing AI model artifacts, training data, and operational records — addressing the residual risk that Frimpong (2026) identifies as 'AI debris.' The SA-CM-MP combination is the only control family grouping that collectively addresses the governance, documentation, and secure disposal obligations of the Retirement lifecycle phase.

Note. Control numbers A-1 through A-6 = Administrative; T-1 through T-5 = Technical; O-1 through O-5 = Operational. Color coding mirrors Appendix A: blue = Administrative, green = Technical, yellow/gold = Operational. SP 800-53 family abbreviations: AC = Access Control; AU = Audit & Accountability; CA = Assessment, Authorization & Monitoring; CM = Configuration Management; IA = Identification & Authentication; IR = Incident Response; MP = Media Protection; PL = Planning; PM = Program Management; RA = Risk Assessment; SA = System & Services Acquisition; SI = System & Information Integrity; SR = Supply Chain Risk Management. NIST AI RMF functions: GOVERN = organizational accountability and policy; MAP = risk identification and categorization; MEASURE = assessment and evaluation; MANAGE = operational enforcement and lifecycle management.