April 2026

Insights

Internal Controls Over Generative AI:

A Uniqus Point of View for Boards, Audit Committees and the C-Suite

Reassessing SOX Compliance, ITGC, Access Controls, and Segregation of Duties in the Age of AI

Informed by COSO's 2026 Guidance: Achieving Effective Internal Control Over Generative AI

EXECUTIVE SUMMARY

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

Generative AI is no longer a future-state consideration for finance and governance leaders — it is already embedded in the processes that produce financial statements, reconcile accounts, extract data from contracts, and generate management analyses. The speed of adoption has outpaced the governance infrastructure at most organizations, creating an urgent need to reassess how internal controls — particularly those mandated under the Sarbanes-Oxley Act (SOX) — address the unique risks introduced by AI-powered systems.

In early 2026, the Committee of Sponsoring Organizations of the Treadway Commission (COSO) published its landmark guidance, "Achieving Effective Internal Control Over Generative AI," which provides the first comprehensive, principle-based framework for integrating GenAI governance into the COSO Internal Control – Integrated Framework (ICIF). This publication introduces a capability-first taxonomy of eight GenAI use case types, maps them to all five COSO components and 17 principles, and provides audit-ready control expectations and metrics.

This Uniqus Point of View distills the implications of the COSO guidance for public companies in the US and their governance leaders. Specifically, we address how existing SOX control frameworks must be reassessed and enhanced, the critical role of IT General Controls (ITGCs), access management, and segregation of duties (SoD) in an AI-enabled environment, practical failure scenarios that illustrate what can go wrong across each GenAI capability type, sector-specific considerations for Technology, Financial Services, and Life Sciences — the three industries at the frontier of AI adoption — and how Uniqus can support organizations through this transition.

Key Takeaway for the Audit Committee and Board

GenAI does not require a separate governance framework — it requires adhering to the established governance basics, with a deliberate extension of your existing COSO-based SOX program. The organizations that act now will convert GenAI from an emerging risk into a well-governed asset; those that delay face regulatory exposure, audit findings, and material weakness risk.

Sharad Chaudhry
Partner, Accounting & Reporting Consulting

Vartika Saxena
Partner, Tech Consulting

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

01

WHY THIS MATTERS NOW: THE CONVERGENCE OF AI ADOPTION AND REGULATORY EXPECTATIONS

The adoption of generative AI in finance functions has accelerated dramatically. Across Corporate America, GenAI tools are being deployed for invoice processing, journal entry preparation, account reconciliation, financial close acceleration, lease abstraction, revenue contract analysis, and drafting management commentary. These are not peripheral experiments — they sit directly in the path of internal control over financial reporting (ICFR).

1.1

The Regulatory Landscape Is Tightening

Several regulatory developments underscore the urgency for governance leaders:

- **PCAOB 2026 Inspection Priorities:** The PCAOB has signaled increased attention to how auditors evaluate the use of AI and automated tools within client environments, including the adequacy of ITGCs over AI systems and the sufficiency of audit evidence where AI-generated outputs are relied upon.
- **SEC Enforcement Posture:** The SEC has made clear that management's responsibility for ICFR extends to any technology used in the financial reporting process, including AI. "AI washing" — overstating AI capabilities in public disclosures — has already drawn enforcement scrutiny.
- **COSO 2026 GenAI Guidance:** The new COSO publication provides an authoritative framework for internal control over GenAI, extending the ICIF principles into AI-specific practices with audit-ready control expectations.
- **EU AI Act:** For US multinationals operating in the EU, the AI Act introduces tiered compliance requirements that intersect with financial reporting processes and internal controls.

The above developments indicate that the overarching control environment needs to evolve and adapt from the current static or periodic update model to a dynamic, more context-based one.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

1.2

The "Shadow AI" Problem

GenAI's low barrier to entry means that finance teams, operational units, and individual employees can adopt AI tools outside formal IT governance channels. This phenomenon — "shadow AI" — is the AI equivalent of shadow IT but with significantly higher risk, because GenAI outputs are probabilistic, can be wrong with high confidence score, and may be used to inform financial reporting decisions without adequate validation. The COSO guidance specifically identifies shadow AI as a critical control environment risk that requires detection mechanisms and clear, acceptable use policies.

1.3

Rate of Change is exponential and irreversible

The pace of technological change is faster than ever, and the acceleration in GenAI development and adoption is unprecedented. Without effective guardrails, human-in-the-loop oversight, and robust governance frameworks, GenAI can quickly become a black box — producing errors or bias that materially and adversely impact financial reporting. These errors or gaps may also be difficult to identify, understand, or reverse if the underlying models are not well-documented and continuously monitored. The COSO guidance addresses both the transformative potential of this technology and its associated risks, providing an effective and adaptable internal control framework that enables organizations to embrace change and create competitive advantage — without compromising reliability or accountability.

02

UNDERSTANDING THE COSO ‘CAPABILITY-FIRST’ TAXONOMY

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO ‘Capability-First’ Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

Unlike prior guidance that organized AI governance by technology type or vendor, the COSO 2026 framework introduces a capability-first taxonomy that classifies GenAI use cases into eight capability types across the data-to-decision lifecycle. This approach is powerful because it focuses on what the AI system actually does, making risk assessment and control design independent of the specific vendor or model used.

#	Capability Type	Financial Reporting Relevance
01	**Data Extraction & Ingestion**	Extracting invoice fields, contract terms, bank statement data; OCR on source documents
02	**Data Transformation & Integration**	Normalizing the chart of accounts, currency conversion, and data mapping between systems
03	**Automated Transaction Processing**	Auto-posting journal entries, matching invoices to POs, and processing intercompany reconciliations
04	**Workflow Orchestration**	AI agents managing financial close tasks, routing approvals, and escalating exceptions
05	**Judgment & Insight Generation**	Impairment analysis, revenue forecasting, lease classification, and going concern assessment support
06	**AI-Powered Monitoring**	Continuous controls monitoring, anomaly detection in journal entries, and fraud pattern identification
07	**Knowledge Retrieval & Summarization**	Summarizing new accounting standards, regulatory updates, and contract analysis for disclosure
08	**Human–AI Collaboration**	Copilots for memo drafting, disclosure review, and technical accounting research assistance

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

Uniqus Perspective:

Mapping Your AI Landscape

Before redesigning controls, every organization must first inventory and classify its GenAI use cases against these eight capability types. In our experience, most organizations have significant GenAI exposure in Capabilities 1–3 (data ingestion, transformation, and transaction processing) but have not yet extended ICFR controls to cover them. Capabilities 4–5 (workflow orchestration and judgment generation) present the highest risk for SOX-relevant processes because they involve autonomous decision-making or influence management estimates.

Establishing Accountability and Responsibility for AI

As a next step in setting up AI governance, establish ownership for AI and responsibility for the outcomes it produces. This ensures a conscious evaluation of the use case, scrutiny of any changes, and review of outcomes. A detailed RACI matrix for all AI assets, datasets, and escalation paths for anomalies and ethical considerations not only makes oversight and governance effective but also ensures that the tone at the top is set and guardrails are applied consistently and effectively, where needed.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

03

WHAT CAN GO WRONG: FAILURE SCENARIOS ACROSS CAPABILITY TYPES

Understanding theoretical risks is important, but governance leaders need concrete illustrations of how GenAI failures materialize in practice. The following section maps a realistic failure scenario for each of the eight capability types to the relevant COSO component(s) that would have prevented or detected the issue. These examples are drawn from patterns observed across our client base and industry research.

3.1

Data Extraction & Ingestion

Scenario: The Phantom Invoice Line Items

A multinational consumer goods company deploys a GenAI extraction tool to read incoming vendor invoices and populate the AP subledger. The system performs well on standard PDF invoices, but encounters scanned fax copies from a legacy supplier in Southeast Asia. The model "hallucinated" a line item for freight charges that did not exist on the original document, inserting a fabricated USD 47,000 charge. Because the amount fell below the auto-posting threshold and matched a typical freight range, it was auto-approved and posted without human review. The error was discovered six weeks later during a quarterly AP aging review.

COSO Linkage

Risk Assessment (Principle 7): The initial risk assessment did not consider document quality variability across suppliers and geographies as a risk factor.

Control Activities (Principle 10): Confidence thresholds were set too high, allowing low-confidence extractions to bypass human review. A risk-calibrated threshold with mandatory human review for scanned/low-quality documents would have caught this.

Monitoring Activities (Principle 16): No ongoing accuracy monitoring by document type. A dashboard tracking extraction accuracy by source format would have flagged the degraded performance on fax documents.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

3.2

Data Transformation & Integration

Scenario: The Silent Currency Mapping Error

A global technology company uses GenAI to normalize and map revenue data from 23 subsidiaries into a consolidated chart of accounts. After a subsidiary in Brazil changed its ERP system, the GenAI transformation logic incorrectly mapped a local revenue sub-category to "Other Income" instead of "Product Revenue." The mapping error was not immediately visible because consolidated totals were unaffected — only the revenue line item disaggregation was wrong. The misclassification persisted for two quarters and was identified only when the external auditor performed detailed revenue disaggregation testing.

COSO Linkage

Risk Assessment (Principle 9): The change management process failed to identify the ERP migration as a significant change requiring transformation logic re-validation.

Control Activities (Principle 10): No automated reconciliation between pre-transformation and post-transformation revenue subtotals by category. A pre/post reconciliation at the sub-category level would have detected the reclassification.

Information & Communication (Principle 14): The subsidiary's ERP migration was not communicated to the central transformation logic owner, so the mapping rules were not updated.

Monitoring Activities (Principle 16): No ongoing monitoring for revenue sub-categories (for anomalies or exceptions against previous values/ patterns). An ongoing monitoring of changes in values/ patterns from previous cycles would have detected the issue earlier.

3.3

Automated Transaction Processing & Reconciliation

Scenario: The Runaway Intercompany Netting

A financial services firm deploys an AI reconciliation agent to match and net intercompany transactions across 40+ legal entities. A misconfigured currency conversion parameter caused the agent to match transactions using stale exchange rates cached from the prior month-end. Over a three-day period, the agent auto-posted 2,300 intercompany entries, resulting in an aggregate USD 12 million overstatement in one entity and a corresponding understatement in another. Because both sides netted to zero at the consolidated level, the error was invisible in the consolidated trial balance and was only discovered during entity-level statutory reporting.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

COSO Linkage

Control Activities (Principle 10): No post-processing analytical review at the entity level. An automated variance check comparing entity-level AI-posted amounts against prior-period or expected ranges would have flagged the anomaly.

Control Activities (Principle 11): The stale exchange rate cache represents an ITGC failure. AI systems that consume reference data must validate data freshness before each processing run.

Monitoring (Principle 16): The error was not detected for three days, indicating insufficient real-time monitoring of auto-posted transaction volumes and amounts.

3.4

Workflow Orchestration & Autonomous Task Execution

Scenario: The Misdirected Close Task

A large industrial company uses an AI orchestration agent to manage its financial close process — pulling trial balance data, assigning reconciliation tasks, escalating exceptions, and compiling close packages. Due to an error in the AI's tagging of "expertise fields," a complex goodwill impairment analysis was automatically routed to a junior staff accountant instead of the senior technical accounting manager. The junior staffer, unfamiliar with ASC 350 requirements, accepted the AI-generated impairment assessment without challenge. The flawed analysis understated an impairment charge by USD 8 million and was caught only during the audit committee review of significant estimates.

COSO Linkage

Control Environment (Principle 4): The AI routing logic failed to match task complexity to reviewer competence. High-risk items, such as impairment analyses, must have routing rules that require a secondary competence check before assignment.

Risk Assessment (Principle 7): Routing misassignment was not identified as a risk scenario during the AI deployment. Scenario analysis should have tested "What if a complex task is routed to an unqualified reviewer?"

Control Activities (Principle 12): The close procedure did not require a documented sign-off confirming that the assigned reviewer had the appropriate expertise for the task's risk level.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

3.5

Judgment & Insight Generation

Scenario: The Hallucinated Lease Classification

A real estate investment trust uses GenAI to analyze new lease agreements and propose classifications under ASC 842. For a complex ground lease with an embedded purchase option, the AI confidently classified the arrangement as an operating lease, citing a present value calculation that appeared reasonable. Upon closer examination during the annual audit, it was discovered that the AI had fabricated the incremental borrowing rate used in its present value calculation — the rate did not correspond to any observable market data or company-specific rate. The misclassification resulted in an understatement of USD 22 million in right-of-use assets and corresponding lease liabilities.

COSO Linkage

Control Environment (Principle 4): The reviewer lacked technical competence in ASC 842 to identify that the borrowing rate was implausible. Reviewers of AI-generated accounting conclusions must have subject matter expertise.

Control Activities (Principle 10): The AI output was treated as a fact rather than an assertion requiring evidence. GenAI outputs in judgment-heavy areas must require source citations that reviewers can independently verify.

Information & Communication (Principle 13): The AI system did not flag that the borrowing rate was self-generated rather than sourced from observable data. Confidence scores and source attribution are essential for judgment outputs.

3.6

AI-Powered Monitoring

Scenario: Alert Fatigue Masks Real Fraud

A retail company implements GenAI-powered monitoring over expense reports and vendor payments to detect fraud patterns. Initially effective, the system generated a high volume of false positives — flagging legitimate international vendor payments as anomalies due to currency fluctuations. Over three months, reviewers developed "alert fatigue" and began dismissing alerts in bulk without individual investigation. During this period, a procurement manager exploited the pattern by creating a fictitious vendor and routing USD 340,000 in payments that were flagged by the AI but dismissed as false positives.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

COSO Linkage

Risk Assessment (Principle 8): The fraud risk assessment did not consider that the AI monitoring system's own weaknesses could become a vector for fraud. Alert fatigue is a predictable failure mode that should be designed against.

Control Activities (Principle 10): No escalation protocol for bulk alert dismissals. Any mass suppression of AI-generated alerts should require supervisory approval and documentation.

Monitoring Activities (Principle 16): The monitoring system itself was not monitored. Precision and recall metrics should have triggered a recalibration when false positive rates exceeded acceptable thresholds.

3.7

Knowledge Retrieval & Summarization

Scenario: The Outdated Regulatory Summary

A multinational pharma company uses GenAI to monitor and summarize evolving data privacy regulations across 30+ jurisdictions for compliance reporting. The system's retrieval-augmented generation (RAG) architecture relies on a regulatory knowledge base that is refreshed monthly. Between refresh cycles, a major jurisdiction enacted significant amendments to its data transfer requirements. The AI continued generating compliance summaries based on the superseded regulation for three weeks, during which the company inadvertently certified compliance with outdated requirements in a regulatory filing.

COSO Linkage

Control Activities (Principle 11): The monthly refresh cycle was insufficient for fast-moving regulatory environments. Critical regulatory sources should have real-time or daily update mechanisms.

Information & Communication (Principle 13): The knowledge base lacked freshness validation. Automated freshness checks with source-level timestamps would have flagged that the regulatory content was stale.

Monitoring (Principle 16): No monitoring of retrieval coverage gaps. A completeness check comparing AI-summarized jurisdictions against a master regulatory watch list would have detected the gap.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

3.8

Human–AI Collaboration

Scenario: The Over-Helpful Disclosure Copilot

A mid-cap software company uses an AI copilot to assist in drafting quarterly earnings release narratives and MD&A sections. During the Q3 preparation cycle, a financial planning analyst used the copilot to draft forward-looking guidance language. The copilot, drawn from internal financial model data accessible through its RAG configuration, included a specific revenue projection range that had not yet been approved by the CEO or discussed with the Board. The draft was circulated to the broader IR team for review. Although caught before external release, the incident exposed the risk of unauthorized disclosure of material non-public information (MNPI) through AI-assisted drafting.

COSO Linkage

Control Environment (Principle 1): The Acceptable Use Policy did not address AI copilot access to MNPI or unapproved forward-looking data. Clear boundaries must define what data sources the copilot can access when drafting external communications.

Control Activities (Principle 12): No pre-release review gate requiring legal/IR sign-off before AI-assisted drafts of market-sensitive documents could be shared beyond the author.

Information & Communication (Principle 15): The incident highlights the need for disclaimers and guardrails in AI interfaces used for external-facing content, including blocks on unapproved financial data.

Uniqus Perspective:

The Common Thread

Across all eight scenarios, failures share a common pattern: organizations treated GenAI outputs with the same level of trust they would give deterministic systems. The fundamental shift required is to treat every GenAI output as a claim requiring validation, not a fact to accept by default. This single mindset change, embedded in control design, reviewer training, and monitoring logic, would have prevented or detected every failure described above.

Another common pattern between all the scenarios above was the lack of a holistic design of control activities. Effective use of GenAI requires a detailed analysis of what could go wrong, identification of edge cases, and building an adequate set of control activities that can support in preventing errors, as well as detective controls that can help identify anomalies and flag for reviews as needed.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

04

REASSESSING THE SOX FRAMEWORK FOR GenAI

The COSO 2026 guidance makes clear that GenAI does not require abandoning or replacing existing SOX programs. Instead, it requires a deliberate reassessment of each COSO component through a GenAI-aware lens. Below, we outline the critical reassessment areas across all five COSO components.

Control Environment: Setting the Tone from the Top

The control environment must explicitly address GenAI governance. Key actions include:

- **GenAI Acceptable Use Policy (AUP):** A board-approved policy defining what data can and cannot be processed by AI, which use cases are permitted in SOX-relevant processes, and boundaries for AI-assisted vs. AI-automated decisions.
- **Cross-Functional AI Governance Committee:** A standing committee comprising legal, compliance, IT, risk management, and finance leadership that reviews GenAI initiatives, assesses risks, and reports to the Audit Committee.
- **Inventory of all GenAI Use Cases:** Mechanisms to be established to maintain a real-time inventory of all GenAI use cases in production and in development, for effective governance and monitoring.
- **Ownership and RACI Matrices:** Every GenAI tool, prompt library, retrieval dataset, and transformation rule must have a named owner with defined authority, escalation paths, and accountability for outcomes.
- **AI Competency Requirements:** Role-based training covering secure prompting, bias recognition, hallucination detection, and GenAI-specific risk factors for control operators, reviewers, and management.

Risk Assessment: Dynamic and AI-Aware

Traditional annual risk assessment cycles are insufficient for GenAI. The COSO guidance emphasizes that risk assessment must be continuous or near continuous because models, prompts, retrieval sources, and vendor configurations can change rapidly. Organizations must:

- Evaluate every GenAI use case for AI-specific threats: hallucinations, bias, prompt injection, model drift, data provenance gaps, and third-party dependencies
- Perform scenario analysis ("What if..." exercises) for each capability type to surface hidden dependencies and edge cases
- Build monitoring or discovery mechanisms to identify GenAI changes in real time and ensure review tollgates are designed (Human in the loop) for critical changes, before going live.
- Maintain living risk registers that update when models, data corpora, or configurations change — not just at annual review cycles
- Assess whether GenAI is the right tool for each use case — in some cases, deterministic automation provides greater reliability at lower risk
- Evaluate novel fraud risks: deepfakes, synthetic records, model manipulation through crafted prompts, and excessive agency granted to AI agents

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

Control Activities: Human-in-the-Loop and Beyond

Control activities must account for the probabilistic nature of GenAI outputs. The COSO guidance identifies five leading approaches for operationalizing controls over AI:

Approach	Description
Human-in-the-Loop (HITL)	Ranges from full re-performance of AI outputs to risk-based sampling of exceptions. Required for high-risk SOX processes.
Performance Testing	Use of test populations to validate AI accuracy and completeness; stress-testing edge cases and boundary conditions.
Multi-Model Validation	Comparison of outputs across independent AI models or deterministic benchmark algorithms to detect inconsistency, bias, or drift.
Data Analytics Monitoring	Continuous monitoring of AI outputs for anomalies or drift, with thresholds calibrated to risk appetite.
Third-Party Validation	Independent review or certification of AI models and outputs, particularly for vendor-supplied AI systems.

AI Reliance Determination:

A critical concept introduced in the COSO guidance is the distinction between ***reliance*** and ***non-reliance*** on AI outputs for ICFR purposes. When management depends on AI-generated outputs as part of the evidence supporting a control's effectiveness, the AI system must meet the same evidence standards expected for ICFR — including documented configurations, sampling rationale, and retained audit evidence.

Information & Communication: Traceability as a Design Principle

GenAI processes must maintain comprehensive audit trails. The COSO guidance requires organizations to capture prompts, inputs, outputs, source references, model/configuration versions, and confidence scores — all of which can affect control conclusions. Key requirements include:

- Source capture and provenance tagging for all data entering AI pipelines
- Centralized prompt libraries and retrieval knowledge sources with role-based access and version control
- Defined communication protocols for incidents, model changes, and known limitations
- External disclosure considerations when GenAI materially affects reporting or compliance

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

Monitoring Activities: Continuous and Multi-Dimensional

Monitoring must be continuous and multi-dimensional, tracking accuracy, precision/recall, coverage, latency, exception volumes, and fairness metrics. Organizations should:

- Deploy real-time dashboards for GenAI KPIs and KRIs alongside traditional control metrics
- Establish explicit triggers for retraining, reconfiguration, or rollback based on monitored thresholds
- Conduct periodic model effectiveness audits using historical and hypothetical data
- Perform adversarial testing exercises to identify vulnerabilities in AI-enabled processes
- Maintain remediation logs with root cause analysis, corrective actions, and follow-up testing results
- Monitoring systems that use GenAI, to ensure detection logic continues to be accurate and relevant

Looking Ahead: Multi-Metric AI Tolerances

The COSO guidance signals that the definition of "control failure" itself may need to evolve for AI systems. Rather than binary pass/fail determinations, organizations will increasingly rely on multi-metric tolerance ranges across dimensions such as task accuracy, data leakage tolerance, bias levels, explainability minimums, and model change velocity. For the control environment to be effective, control assessments themselves must be continuous, dynamic, and contextual. Equally important, the competency of personnel overseeing control and monitoring activities must be critically evaluated and actively developed.

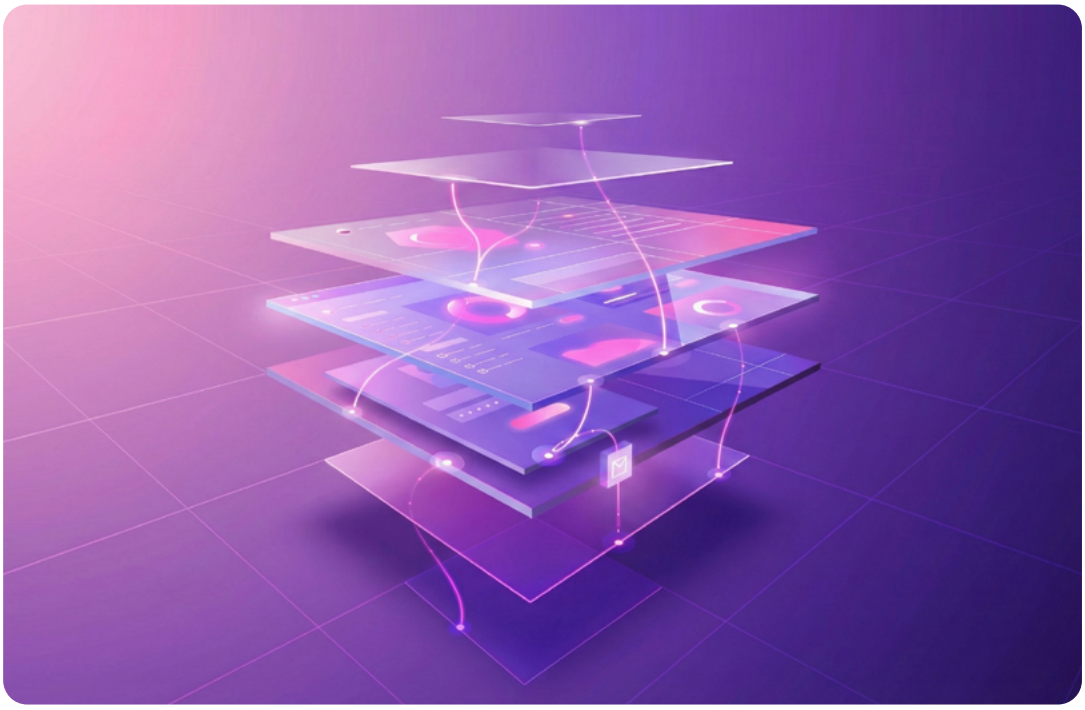

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

05

IT GENERAL CONTROLS (ITGCs) IN AN AI-ENABLED ENVIRONMENT

ITGCs have long been the foundation of SOX compliance for technology environments. The introduction of GenAI requires a fundamental expansion of the ITGC scope to address AI-specific technology components.

Expanding the ITGC Perimeter

Traditional ITGC Scope	GenAI Extension	Control Implications
Application configurations	Prompts, system prompts, temperature settings, and retrieval connectors	Version control, change approval, and rollback capability
Database integrity	RAG indexes, vector databases, fine-tuning artifacts, and embeddings	Access control, data lineage, and freshness validation
Change management	Model updates, vendor-pushed changes, prompt revisions, threshold adjustments	Pre-deployment testing, documented approvals, and post-change monitoring
Batch job scheduling	AI agent orchestration, multi-step autonomous workflows	Simulation testing, routing logic documentation, SoD validation
Interface controls	API integrations with AI models, plugin configurations	Input/output validation, rate limiting, error handling

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

Change Management for AI Systems

Change management is perhaps the most critical ITGC extension for GenAI. Unlike traditional applications, where changes are discrete and planned, GenAI systems can change through multiple vectors:

- **Vendor-initiated model updates:** Cloud-based AI providers may update models, safety filters, or capabilities without explicit customer approval. Organizations must negotiate contractual notification obligations and implement independent verification procedures.
- **Prompt and configuration drift:** In collaborative environments, prompts and system configurations may be modified by multiple users without formal change control. All prompts in SOX-relevant processes must be treated as governed configuration items.
- **Data corpus evolution:** For RAG-based systems, changes to the underlying knowledge base directly affect AI behavior. Document freshness, completeness, and accuracy must be continuously validated.
- **Model drift:** Performance degradation over time as data distributions shift. Continuous monitoring with defined retraining triggers is essential.

Logical Access Controls and Identity Management

GenAI introduces new dimensions to access control that extend beyond traditional user provisioning:

- **Model Access:** Who can invoke AI models, modify system prompts, adjust confidence thresholds, and configure retrieval sources? These permissions must be defined, provisioned, and periodically recertified.
- **Data Access through AI:** GenAI systems can access and aggregate data across multiple sources in ways that bypass traditional access restrictions. A user who lacks direct database access may extract sensitive information through AI queries.
- **AI Agent Permissions:** Autonomous AI agents must operate under the principle of least privilege, with clearly defined permission boundaries, logging, and human override capabilities.
- **API Key and Token Management:** API keys connecting internal systems to external AI providers must be managed, rotated, and monitored with the same rigor as privileged credentials.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

06

SEGREGATION OF DUTIES (SoD) IN AN AI-DRIVEN ENVIRONMENT

Segregation of duties is a foundational control principle profoundly challenged by GenAI. When AI systems can simultaneously configure processing rules, execute transactions, and generate review reports, traditional SoD frameworks require fundamental re-examination.

New SoD Conflicts Introduced by GenAI

SoD Conflict	Traditional Control	GenAI Challenge	Recommended Mitigation
Configure vs. Execute	Different roles configure and execute transactions	A single AI system may configure rules and execute simultaneously	Separate prompt/ config access from production execution; require multi-party approval for changes
Execute vs. Review	Maker-checker model with human segregation	AI generates output and may generate review documentation, creating circular validation	Independent human validation of AI review outputs; prohibit AI self-review in SOX processes
Develop vs. Deploy	Developers cannot promote to production	Prompt engineers may modify production prompts without deployment gates	Formal CI/CD pipeline for prompt changes with independent testing and approval gates
Authorize vs. Record	Authorization and recording by different individuals	AI agents may approve, record, and reconcile within a single workflow	Decompose agent permissions; separate authorization from recording at the system level

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

07

ASSESSING FRAUD RISK IN AN AI-ENABLED PROCESS

Fraud risk is heightened in varied ways when there is the use of AI-enabled workflows, task orchestrations, and monitoring. While having a human in the loop can mitigate this risk to some extent, a more 'fraud-conscious' mindset is required to effectively use these capabilities. Autonomous agents can combine multiple, traditionally separate financial functions into a single workflow or a single identity. Beyond the cases of SoD conflicts discussed above, there are a few more scenarios that require careful consideration.

Authorization Risks

Examples

Single AI agents having excessive permission across multiple systems or modules, not originally scoped, could create silent SoD conflicts, leading to risk of fraudulent transactions.

Recommended Mitigation

Create AI agents as non-human identities and apply strict, well-defined, role-based access controls. Ensure periodic monitoring of the non-human identities and set up alerts for any critical access updates.

Excessive Reliance

Examples

An agent with permissions to remediate exceptions might adjust transaction amounts or alter account codes to pass automated checks.

An agent might present a recommendation after a complex analysis to a human in the loop who just approves it without revalidation.

Recommended Mitigation

Create effective monitoring controls for remediation activities and confidence scores for the approach/corrections recommended by the agents. A sample of these must be mandatorily reviewed by a human to maintain confidence in the actions taken by AI agents.

Insecure Interfaces leading to Collusion with External Actors

Examples

An AI agent could be manipulated through prompt injections or by a malicious actor, leading it to act on fraudulent instructions while simultaneously presenting a legitimate audit log for human review, making it difficult to identify such fraudulent transactions.

Recommended Mitigation

Periodic vulnerability assessments need to be performed to identify risks to the AI agent.

Toxic combinations should be avoided when the agent has access to sensitive data, exposure to untrusted content, and the ability to communicate externally. Real-time monitoring of external communications, with alert mechanisms, needs to be in place.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

08 SECTOR-SPECIFIC CONSIDERATIONS

While the COSO framework and the control principles discussed in this POV apply universally, the pace of GenAI adoption, the nature of AI use cases, and the regulatory overlay vary significantly by industry. Technology companies are leading adopters, but Financial Services and Life Sciences / Pharmaceuticals are not far behind — and in some cases, face even more stringent regulatory requirements. Below, we examine the unique GenAI control challenges and priorities for each sector.

8.1 Technology Sector: The First Movers

Technology companies are the natural early adopters of GenAI in their finance functions, given their organizational comfort with AI, engineering talent density, and the competitive imperative to demonstrate AI-native operations. However, this very speed of adoption creates governance gaps that are now drawing regulatory attention.

Unique GenAI Use Cases in Tech Finance

- Revenue recognition automation for complex SaaS arrangements under ASC 606 — AI-assisted identification of performance obligations, standalone selling prices, and variable consideration in multi-element contracts
- Stock-based compensation modeling under ASC 718 with AI-generated valuation assumptions and forfeiture estimates
- Capitalization vs. expense determination for internal-use software under ASC 350-40, where AI classifies development activities across project stages
- AI copilots embedded in financial planning and analysis (FP&A) tools that generate guidance ranges, scenario analyses, and investor-facing narratives

Key Risk Factors

- **Shadow AI prevalence:** Tech employees are more likely to adopt AI tools independently, creating widespread shadow AI that bypasses governance. Engineering teams may build internal AI tools that process financial data without the awareness of finance or compliance teams.
- **Vendor-as-customer complexity:** Many tech companies are simultaneously AI vendors and AI consumers, creating unique SoD and independence concerns — particularly when the company's own AI products are used in its finance function.
- **Rapid iteration culture:** The "move fast" ethos can conflict with the deliberate, documented change management required for SOX-relevant processes. Prompt changes may be deployed to production without formal approval or testing.
- **DISE implications:** The new Disaggregation of Income Statement Expenses (ASC 220-40) requirements amplify the risk of AI misclassification errors in expense categorization, as granular functional disaggregation demands higher precision than consolidated reporting.

Uniqus Recommendation for Tech Companies

Establish a mandatory AI registry that captures every GenAI instance touching financial data, including engineering-built tools. Implement "AI gates" in the SDLC that require finance and compliance sign-off before any AI capability that touches financial reporting data moves to production. Prioritize SoD assessment for cases where the company uses its own AI products in its finance function.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

8.2

Financial Services: Regulatory Intensity Meets AI Ambition

Banks, insurance companies, asset managers, and fintech firms are aggressively deploying GenAI across trading, credit decisioning, claims processing, regulatory reporting, and financial close operations. The sector's unique regulatory overlay — including OCC, Federal Reserve, FDIC, SEC, and state insurance regulators — creates an especially complex governance environment for AI adoption.

Unique GenAI Use Cases in Financial Services

- Credit loss estimation under ASC 326 (CECL) with AI-generated probability of default (PD), loss given default (LGD), and exposure at default (EAD) models incorporating macroeconomic scenario generation
- Insurance claims processing and reserve estimation using GenAI to analyze policy language, extract claim details, and propose IBNR reserve adjustments under ASC 944
- Regulatory reporting automation (Call Reports, FR Y-9C, FFIEC) where GenAI maps transaction data to complex regulatory line items across multiple jurisdictions
- Anti-money laundering (AML) and sanctions screening, where GenAI generates suspicious activity narratives and prioritizes alerts for investigation
- Fair value measurement under ASC 820 with AI-assisted valuation of Level 2 and Level 3 instruments using market data synthesis

Key Risk Factors

- **Model risk management (SR 11-7):** The Federal Reserve's SR 11-7 guidance on model risk management applies to AI models used in financial reporting. Financial institutions must determine whether GenAI systems meet the definition of a "model" and, if so, subject them to the full model validation lifecycle — creating a dual governance requirement alongside SOX.
- **Materiality amplification:** Given the scale of financial services balance sheets, even small AI errors in reserve estimation, fair value measurement, or regulatory reporting can produce material misstatements measured in hundreds of millions of dollars.
- **Explainability requirements:** Regulators and auditors expect transparency in how AI-influenced estimates are derived. GenAI's opaque reasoning frustrates the explainability expectations that are foundational to regulatory examinations.
- **Third-party vendor concentration:** Heavy reliance on a small number of cloud-based AI providers creates systemic concentration risk that banking regulators are increasingly scrutinizing.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

Uniqus Recommendation for Financial Services

Integrate AI governance with existing model risk management (MRM) frameworks under SR 11-7 to avoid duplicative governance structures. Perform a materiality-driven inventory of GenAI use cases, prioritizing CECL, fair value, and regulatory reporting processes. Develop "explainability standards" that define the minimum level of reasoning transparency required for each use case tier. Engage the external auditor early on expectations for AI-related audit evidence in highly regulated financial reporting areas.

8.3

Life Sciences & Pharmaceuticals: Data Integrity Meets Innovation

Life sciences and pharmaceutical companies are rapidly deploying GenAI across clinical data management, regulatory submission preparation, pharmacovigilance, revenue operations, and financial reporting. The sector's unique intersection of FDA/EMA regulatory requirements, complex revenue arrangements, and massive R&D capitalization decisions creates distinctive GenAI governance challenges.

Unique GenAI Use Cases in Life Sciences

- Revenue recognition under ASC 606 for complex collaboration and licensing arrangements, where GenAI analyzes multi-element contracts to identify performance obligations, variable consideration (milestones, royalties), and constraint estimates
- R&D capitalization analysis where GenAI classifies research activities as pre-approval (expensed) or post-approval (capitalized), directly affecting reported earnings
- Inventory valuation and write-down estimation for pharmaceutical products with shelf-life constraints, where GenAI analyzes demand forecasts, regulatory status, and expiration data under ASC 330
- Contract data extraction from complex licensing, co-development, and co-promotion agreements where AI identifies embedded derivatives, variable payment triggers, and accounting-relevant terms
- Pharmacovigilance and litigation reserve analysis where GenAI summarizes adverse event data and legal exposure to support contingent liability assessment under ASC 450

Key Risk Factors

- **FDA 21 CFR Part 11 intersection:** When GenAI processes or generates data that feeds into systems subject to FDA electronic records requirements, the AI system's audit trail, access controls, and validation requirements must satisfy both SOX and FDA standards simultaneously.
- **Revenue recognition judgment complexity:** Pharma revenue arrangements are among the most judgment-intensive under ASC 606. AI-assisted analysis of performance obligations and variable consideration creates high hallucination risk in areas where the financial impact can be in the billions.

- **Data integrity culture:** The life sciences industry has a deeply ingrained data integrity culture driven by FDA requirements (ALCOA+ principles). Extending this discipline to GenAI outputs requires adapting existing data integrity frameworks to address the probabilistic nature of AI.
- **Multi-jurisdictional regulatory complexity:** Global pharma companies must navigate EU AI Act requirements (high-risk AI system classification for medical applications), FDA guidance on AI/ML, and local regulatory expectations across 50+ markets, creating a patchwork of compliance obligations.

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

Uniqus Recommendation for Life Sciences

Leverage existing data integrity frameworks (ALCOA+) as a foundation for GenAI control design, extending attributable, legible, contemporaneous, original, and accurate requirements to AI-generated outputs. Prioritize control design for revenue recognition and R&D capitalization use cases, where AI hallucination risk intersects with high-judgment, high-materiality estimates. Conduct a gap analysis against both SOX and FDA 21 CFR Part 11 requirements for any GenAI system that touches data flowing into regulated environments.

Cross-Sector Observation

Regardless of sector, the organizations making the most progress on GenAI governance share three characteristics: (1) they have established a centralized AI inventory before scaling adoption, (2) they treat prompts and AI configurations as governed assets with the same rigor as application code, and (3) they have engaged their external auditors proactively on AI control expectations rather than waiting for audit findings.

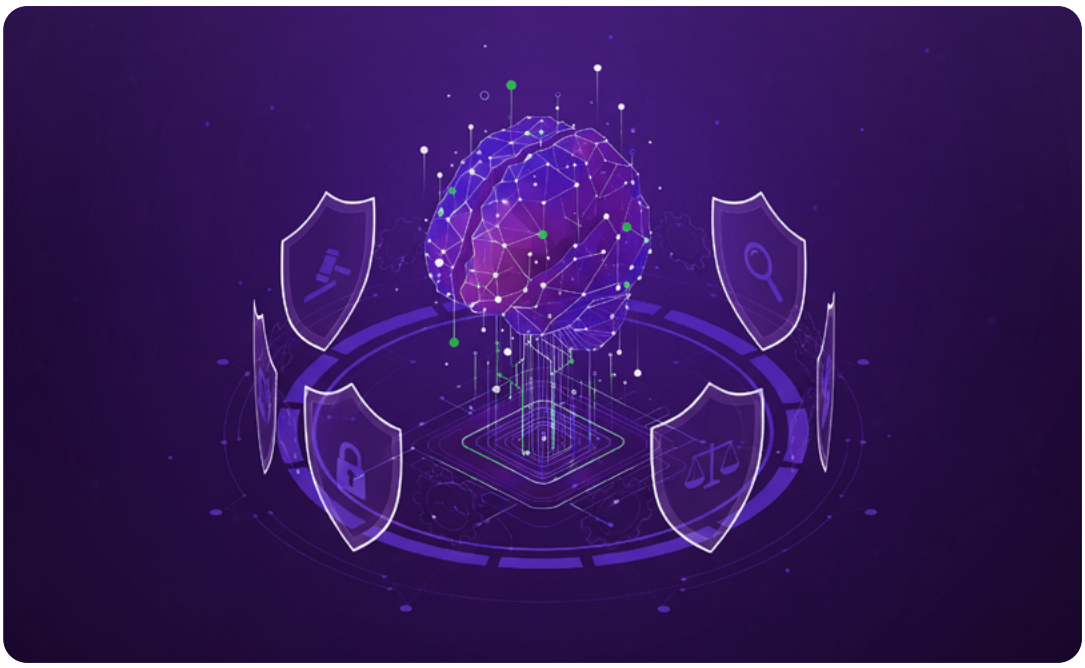

09

IMPLEMENTATION ROADMAP: A SIX-STEP APPROACH

The COSO guidance provides a practical, cyclical implementation roadmap. We have adapted this into a Uniqus-informed approach that reflects our experience advising US public listed companies on SOX compliance transformation:

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

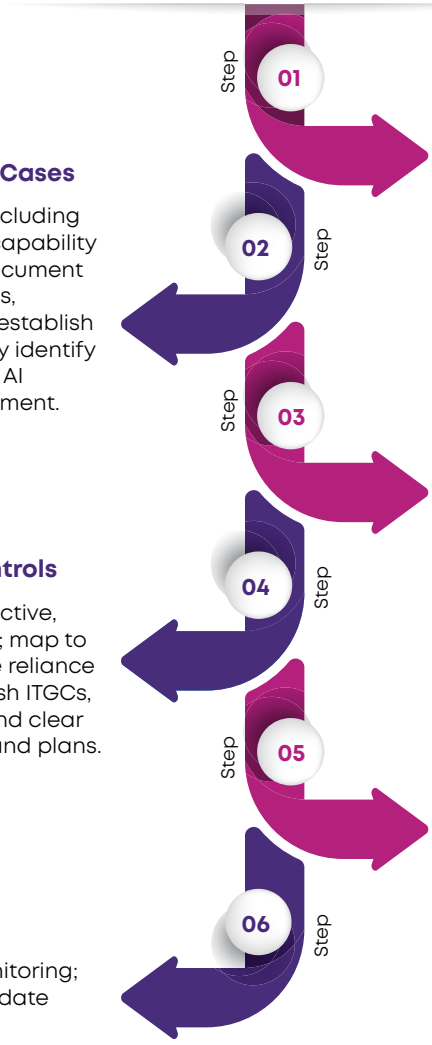

Establish AI Governance

Form a cross-functional AI committee; define charter and Audit Committee interface; establish AI risk appetite statement, conduct role-specific AI trainings to ensure controls are supported by AI-aware humans in the loop, where necessary

Inventory GenAI Use Cases

Comprehensive scan including shadow AI; classify by capability type; assign owners; document objectives, data sources, models, and criticality; establish processes to continually identify the introduction of new AI models into the environment.

Assess Risks

Evaluate each use case against all five COSO components; assess AI-specific risks; perform scenario analysis; build living risk registers; build feedback mechanisms from control monitoring to risk assessment, to ensure any anomalies detected inform the risk assessment process for identification of an adequate level of required controls.

Design and Map Controls

Design preventive, detective, and corrective controls; map to COSO principles; define reliance determinations; establish ITGCs, SoD, access controls, and clear rollback requirements and plans.

Implement & Communicate

Deploy controls; configure dashboards and alerts; train operators and reviewers; publish escalation paths and change protocols.

Monitor and Adapt

Continuous KPI/KRI monitoring; drift reviews; vendor update validation; governance reporting; model effectiveness audits; return to Step 1

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

10

WHAT THE AUDIT COMMITTEE AND BOARD SHOULD ASK

The most effective governance starts with the right questions. Below are the critical questions Audit Committee members and the Board of Directors should be asking management:

Inventory and Awareness

- Do we have a complete inventory of GenAI tools and use cases, including shadow AI?
- Which GenAI use cases directly affect financial reporting processes or SOX-relevant controls?
- Has management classified each use case by capability type and assigned named owners?

Governance and Policy

- Do we have a board-approved GenAI Acceptable Use Policy addressing data classification, prohibited uses, and ethical boundaries?
- Is there a cross-functional AI governance committee with regular Audit Committee reporting?
- Are AI competency requirements defined for control operators, reviewers, and management?

Risk Assessment and Controls

- Has the ICFR risk assessment been updated to include AI-specific risks such as hallucinations, bias, model drift, and prompt injection?
- For each SOX-relevant AI use case, has management determined whether it relies on AI outputs for control effectiveness?
- Are ITGCs extended to cover AI models, prompts, retrieval sources, and vendor configurations?
- Have SoD matrices been updated to reflect AI system roles and permissions?

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

Sector-Specific Considerations

- Have we assessed the intersection of GenAI governance with sector-specific regulatory requirements (e.g., SR 11-7 for banking, 21 CFR Part 11 for pharma)?
- Are our AI controls designed to address sector-specific use cases and risk factors?

Monitoring and Audit Readiness

- Are continuous monitoring mechanisms in place with defined triggers for retraining or rollback?
- Is sufficient evidence being retained to support audit conclusions for AI-enabled controls?
- Has the external auditor been engaged in their approach to evaluating AI-related controls?

Executive Summary

Why This Matters Now: The Convergence of AI Adoption and Regulatory Expectations

Understanding the COSO 'Capability-First' Taxonomy

What Can Go Wrong: Failure Scenarios Across Capability Types

Reassessing the SOX Framework for GenAI

IT General Controls (ITGCs) in an AI-Enabled Environment

Segregation of Duties (SoD) in an AI-Driven Environment

Assessing Fraud Risk in an AI-enabled process

Sector-Specific Considerations

Implementation Roadmap: A Six-Step Approach

What the Audit Committee and Board Should Ask

How Uniqus Can Help

HOW UNIQUS CAN HELP

Uniqus Consultech is uniquely positioned to help organizations navigate the intersection of GenAI adoption and SOX compliance. As a global consulting platform that does not provide audit or tax services, we offer independent advisory without the conflict-of-interest constraints that limit traditional Big 4 firms.

GenAI Control Design & SOX Integration

- AI use case discovery, inventory, and capability-type classification
- AI-specific risk assessment aligned to COSO 2026
- Control design for all eight capability types
- Reliance determination framework for ICFR

ITGC & Access Control Enhancement

- ITGC perimeter expansion for AI systems
- Change management framework for AI configurations
- Access control redesign, including AI agent permissions
- SoD matrix development for AI-enabled processes

Sector-Specific AI Governance

- Tech sector: Shadow AI detection and SDLC AI gates
- Financial services: SR 11-7 and AI integration
- Life sciences: FDA 21 CFR Part 11 gap analysis
- Cross-sector AI governance maturity assessment

Audit Committee & Board Advisory

- Board-ready AI governance briefing materials
- GenAI Acceptable Use Policy development
- Regulatory landscape briefings (SEC, PCAOB, EU AI Act)
- Continuous monitoring dashboard design

About Uniqus Consultech:

Uniqus Consultech is a global tech-enabled consulting company that specializes in Accounting & Reporting, Governance, Risk & Compliance, Sustainability & Climate, Tech Consulting, and Valuations. The Company is co-founded by consulting veterans Jamil Khatri and Sandip Khetan and backed by marquee investors such as Nexus Venture Partners, Sorin Investments, and UST.

Uniqus has a global team of 700+ professionals led by 85+ Partners & Directors across eleven offices in the USA, the Middle East, and India. The company serves more than 300 clients, including marquee names in each of the markets it operates in.

Uniqus is committed to leveraging technology and an integrated global delivery model to provide best-in-class consulting services to its clients.

About Our Accounting & Reporting Consulting Services:

In today's dynamic business environment, organizations are continuously required to adapt to the dynamic shifts in accounting & reporting standards and regulations. Enhanced management reporting requirements and accelerated reporting timelines necessitate increased automation within the finance function. Special events, such as preparing for an IPO/capital market transaction or fulfilling the reporting requirements arising from an M&A transaction, introduce additional complexities and challenges for the finance function.

As a result, organizations need to transform their finance functions from time to time with a focus on people, processes, and technology. Our Accounting & Reporting Consulting (ARC) practice is designed to partner with the finance function through these challenges. Our teams function as an extension of your finance teams to execute your plans seamlessly.

We bring the necessary technical expertise, skills, technology, and bandwidth, enabling you to navigate your plans in real time without adding additional manpower costs within the organization.

A Team That You Can Trust To Deliver

Jamil Khatri
Co-Founder & CEO

Sandip Khetan
Co-founder & Global Head of Accounting & Reporting Consulting

Anthony Vitale
Partner, Accounting & Reporting Consulting

Arda Kaya
Partner, Accounting & Reporting Consulting

Avinash Ramkumar
Partner, Accounting & Reporting Consulting

Sharad Chaudhry
Partner, Accounting & Reporting Consulting

Vartika Saxena
Partner, Tech Consulting

Harsimran Singh Sangha
Managing Director, Accounting & Reporting Consulting

Shoichi Ohno
Managing Director, Accounting & Reporting Consulting

Scan this code to find
more content like this

Visit us at *www.uniqus.com*

or follow us on