WHITE PAPER NO. 1 | VERSION 2.0

Enterprise Execution Control in Autonomous Environments

Infrastructure Governance Architecture

Original Publication: April 2026

Current Version: July 2026

xist

Table of Contents

Executive Summary

Enterprise control architecture was built for a world where people initiated consequential action. That world is ending.

Artificial intelligence introduces autonomous execution. Systems can now initiate actions at speeds traditional oversight was not designed to govern. Existing control models are not adequate for this shift.

The governance problem is structural. Legacy controls assume operating models built around human-initiated action. Artificial intelligence breaks that foundational premise. Autonomous execution requires governance infrastructure purpose-built to manage machine-initiated actions, not controls adapted from outdated business models.

The Infrastructure Governance Architecture (IGA) presented in this paper addresses that gap. IGA is a control architecture built for automated execution. IGA separates identity, authorization, and enforcement into an infrastructure model designed for autonomous environments.

The architecture repositions human direction earlier in the execution lifecycle by converting governance decisions into infrastructure-level authority rules. The infrastructure then resolves and enforces those rules before autonomous action occurs.

As AI becomes embedded in enterprise operations, organizations will need infrastructure built specifically to govern execution with reliability, accountability, and control.

IGA delivers that infrastructure.

I. Blockchain as a Foundation for Governed Execution Infrastructure

Blockchain technology was introduced in 2008 as a decentralized ledger for peer-to-peer value transfer.[1] Since its inception, this technology has evolved considerably. Public blockchain networks, particularly Ethereum, are now incorporated into regulated financial products operating under established regulatory frameworks.[2]

These developments demonstrate that blockchain has moved beyond experimental transaction recording. The technology now provides the architectural properties that machine-mediated governance requires: cryptographic identity supports accountable actors, append-only records support auditability and traceability, and programmable logic supports consistent rule enforcement before execution.[3]

The Infrastructure Governance Architecture (IGA) builds on these blockchain-based capabilities as the structural foundation for governed execution.

[1] Blockchain is a distributed ledger technology that records transactions across a network of nodes in an append-only, cryptographically linked chain of blocks. It is designed to operate without institutional intermediaries or a central authority controlling the ledger. Each block is cryptographically linked to the one before it, creating a tamper-evident record in which confirmed entries cannot be altered without detection. The distributed structure of the network also prevents any single party from unilaterally rewriting the ledger's history.

[2] Examples include tokenized fund structures, blockchain-based settlement initiatives by major financial institutions, regulated custody frameworks for digital assets, and other supervised market activities incorporating public blockchain infrastructure. These developments indicate that blockchain networks, including Ethereum, are increasingly being used as operational infrastructure within regulated financial market activity.

[3] Smart contracts are self-executing programs written in code that enforce predefined terms automatically when stated conditions are satisfied. They operate deterministically and do not require discretionary human intervention during execution. Their use in regulated and supervised financial environments is established in practice. Illustrative examples include Franklin Templeton's blockchain-enabled U.S. registered money market fund, J.P. Morgan's live tokenized collateral workflow involving BlackRock and Barclays, and the Swiss National Bank's tokenized-asset settlement pilot on SIX Digital Exchange. Together, these examples demonstrate that smart contracts are already functioning as programmable financial infrastructure, at institutional scale, within regulated markets.

II. Infrastructure Governance Architecture

Autonomous systems operate at machine speed. Traditional governance mechanisms do not.[4] The IGA is designed to close that gap by unifying and relocating control earlier in the execution lifecycle.

To accomplish this, IGA separates execution authority into three independent control domains:

- Identity Layer
- Authorization Layer
- Enforcement Layer

The domains are architecturally separated to preserve effective control across machine-mediated environments. Identity is established independently. Authority is determined dynamically against the applicable conditions for the requested action. Enforcement begins only after that determination permits execution.

Blockchain-based infrastructure is well suited to this model. It provides persistent identity state, deterministic rule execution, and tamper-evident records. These capabilities allow the control logic of IGA to be implemented directly within enterprise infrastructure.[5]

The IGA, with blockchain-based infrastructure, enables unified governance to operate at the pace of automation.[6]

[4] Existing control frameworks address important aspects of identity, access, and policy enforcement. Zero Trust governs trust validation and access context. Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC) govern permission assignment and evaluation based on roles and defined attributes. Policy-as-Code formalizes rule expression and evaluation, and workload identity systems establish verifiable machine identity. Taken together, these functions support important control objectives, yet they are typically implemented across separate systems, applications, and embedded logic rather than unified through a single governed control point for consequential execution.

[5] IGA is technology-agnostic. This paper emphasizes a blockchain-native implementation path, as blockchain satisfies the technological requirements of the architecture: verifiable identity state, governed rule evaluation, and tamper-evident records. A future technology that satisfies those same requirements could also support IGA.

[6] IGA applies to consequential, repeatable actions with defined execution pathways that can be expressed in governed rule form before execution. IGA does not apply to actions requiring open-ended human judgment that cannot be fully specified in advance.

III. Identity Layer

The Identity Layer is the first of three distinct layers within the IGA and establishes the identity artifact upon which subsequent authorization and enforcement decisions depend. Identity within the IGA is defined by two architectural properties: role-specific definition and artifact immutability.

Role-Specific Definition

Role-specific definition is the formal establishment of identity for a declared institutional role or authority context. In the IGA, the identity artifact identifies the capacity in which an actor participates in governed execution, not merely the existence of the actor. Role-specific definition ensures that authorization is evaluated within the correct institutional context.

Artifact Immutability

Artifact immutability is the fixed and verifiable state of an identity artifact after issuance. In the IGA, artifact immutability is maintained by preventing alteration after establishment. Artifact immutability preserves control integrity, traceability, and evidentiary value. The reliability of the control environment depends on the identity artifact remaining fixed after issuance.

Role-specific definition and artifact immutability establish the Identity Layer as the foundational control layer for governed machine execution.

IV. Authorization Layer

The Authorization Layer is the second of three distinct layers within the IGA and evaluates the identity artifact established by the Identity Layer against the rules governing that specific action. Authorization within the IGA is defined by two architectural properties: rule evaluation and authorization decision.

Rule Evaluation

Rule evaluation is the formal assessment confirming the conditions for authority have been satisfied. In the IGA, the identity artifact is evaluated against governing logic and applicable institutional context. Rule evaluation grounds authorization in explicit institutional logic and prevents execution from proceeding on assumption, inference, or incomplete context.

Authorization Decision

An authorization decision is the formal output of rule evaluation. It determines whether the requested action is authorized for execution under the applicable role, rules, and context. In the IGA, the authorization decision is the control point that permits or prevents enforcement. It converts evaluated authority into an executable determination.

Rule evaluation and authorization decision establish the Authorization Layer as the control point for governed machine execution.

V. Enforcement Layer

The Enforcement Layer is the third of three distinct layers within the IGA and validates the authorization output issued by the Authorization Layer. Enforcement is the point at which an authorized determination is converted into controlled system action. Enforcement within the IGA is defined by two architectural properties: conditional execution and system integration.

Conditional Execution

Conditional execution is the controlled initiation of action based on a valid authorization output. In the IGA, enforcement relies on the authorization decision and initiates action only when that decision permits execution. This preserves the separation between authorization and enforcement while ensuring execution occurs only from an authorized determination.

System Integration

System integration is the controlled interface with the downstream system required to carry out an authorized action. In the IGA, system integration is applied by identifying and initiating the specific interface responsible for execution. This ensures that authorized execution is routed to the correct downstream system in a controlled and traceable manner.

Conditional execution and system integration establish the Enforcement Layer as the bridge between authorization and execution.

VI. Architecture Overview

Figure 1 illustrates the three-layer IGA described in the preceding sections. Each layer operates as an independent control domain with explicit structural separation, enabling governed identity establishment, deterministic authorization evaluation, and controlled enforcement across systems.

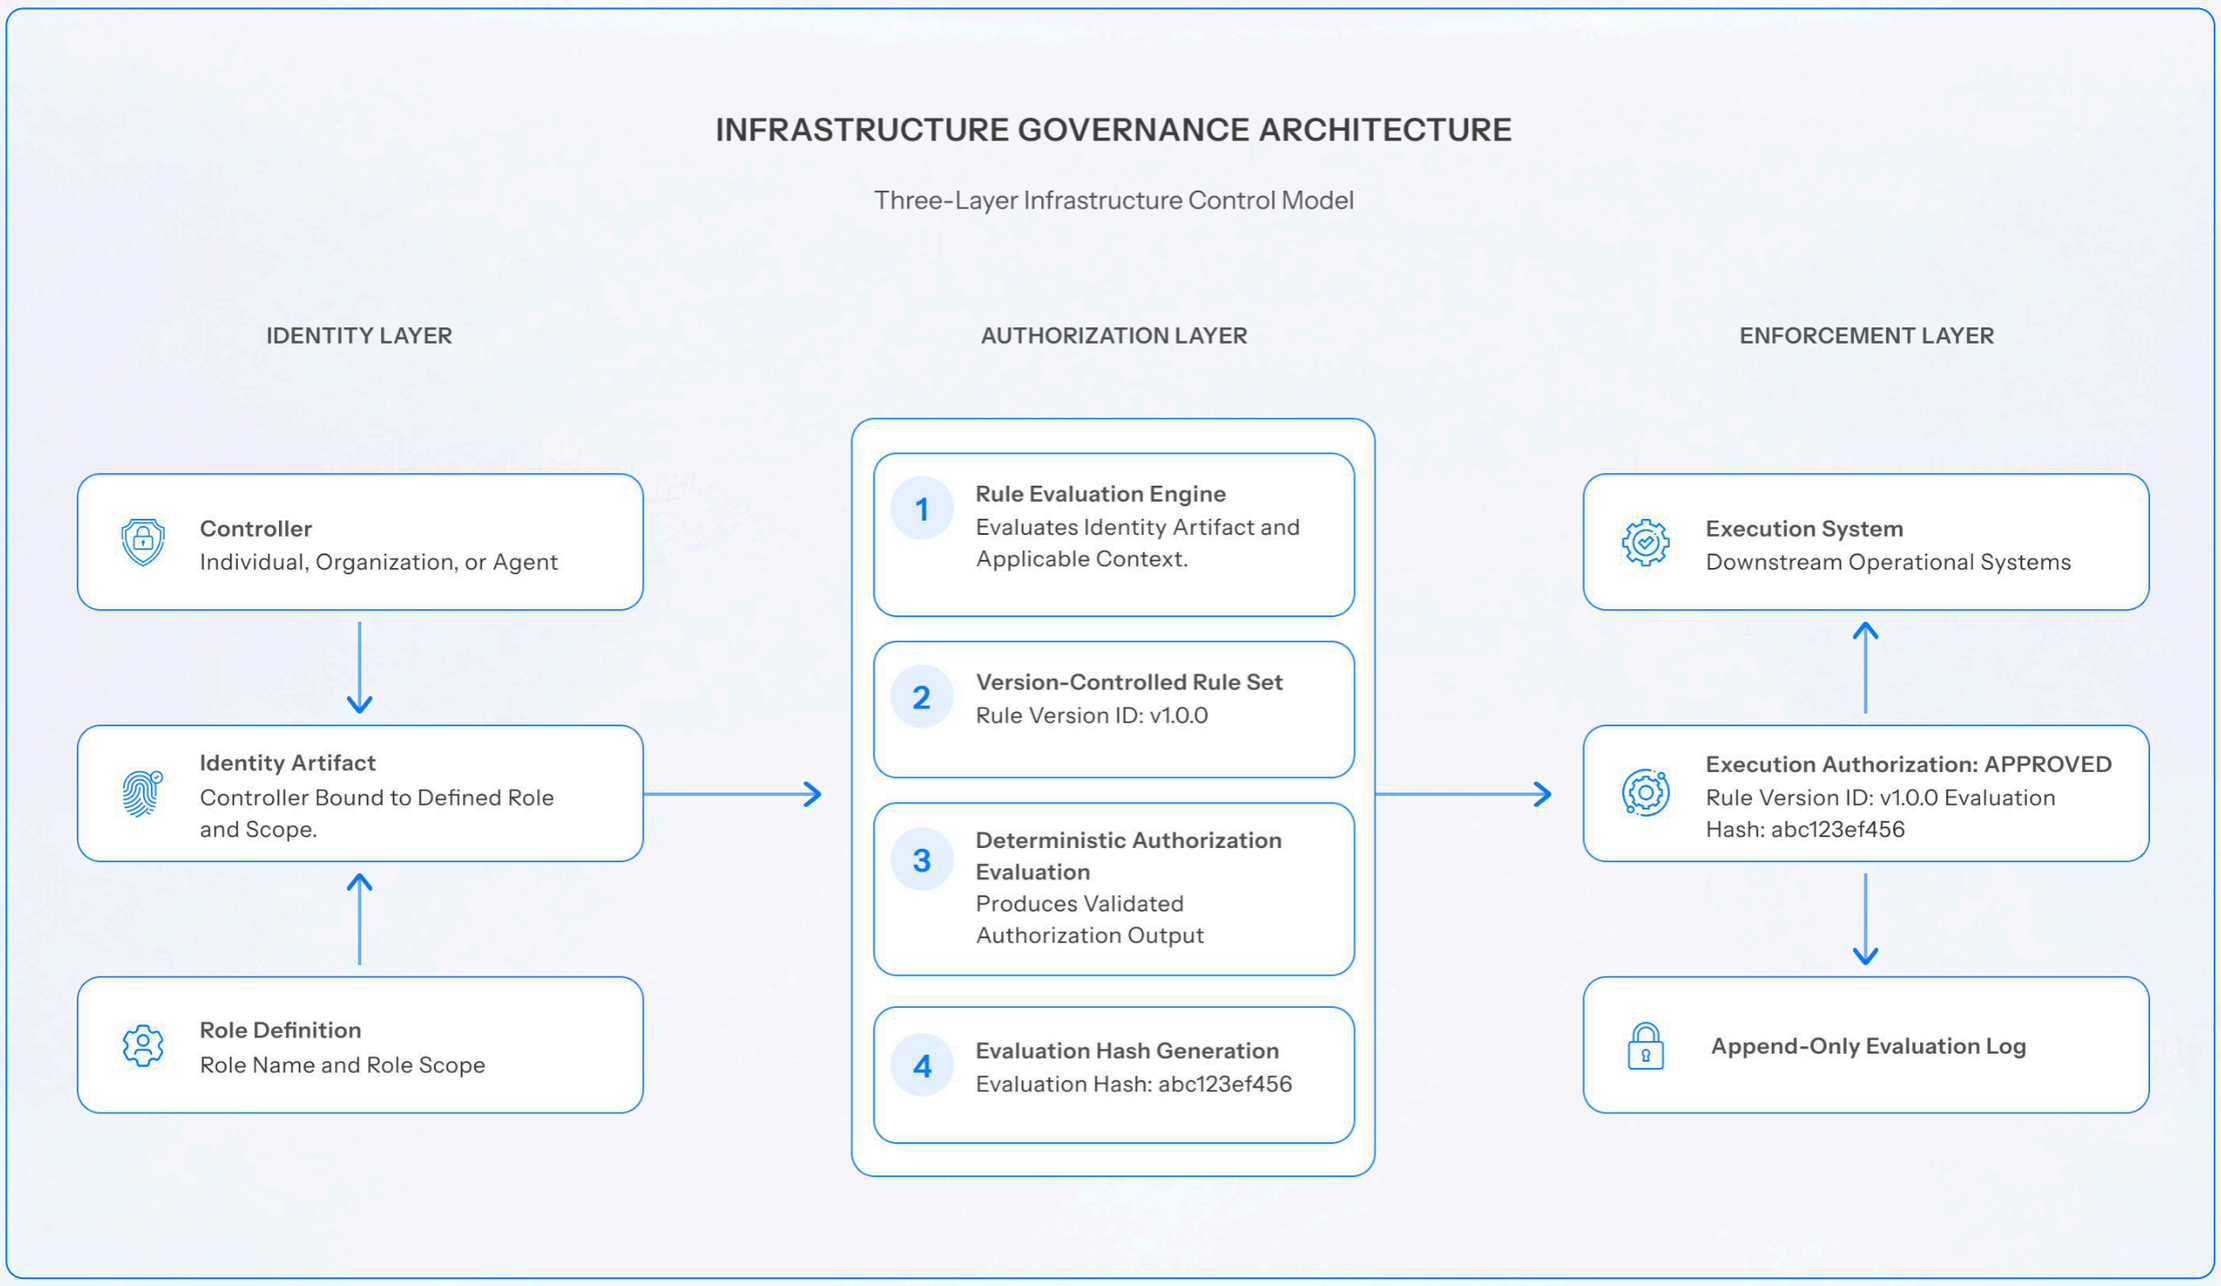

Figure 1. Infrastructure Governance Architecture, Three-Layer Control Model

The model separates identity establishment, authorization evaluation, and operational enforcement into distinct control functions, preserving control integrity, traceability, and governed execution across distributed and autonomous environments. It further depicts illustrative implementation elements that support those domains without restating the architectural attributes defined in Sections III through V.

VII. Enterprise Governance Implications

Effective governance of automated execution requires enterprise-level oversight. As AI systems begin to initiate actions, affect outcomes, and operate across business functions, organizations can no longer treat execution control as a narrow technology issue.

Traditional control domains address financial reporting, compliance, and operational risk. They also address information security and system access. Automated execution cuts across those domains, yet none provides complete governance over machine-directed action.

Enterprises therefore need a defined governance model for automated execution. Boards and senior management require clear accountability, transparent oversight, and enforceable control over systems acting on behalf of the organization.

IGA provides that governance model. It establishes automated execution as an enterprise control responsibility and gives organizations the infrastructure to manage machine-directed action with reliability, accountability, and control.

VIII. Conclusion

Autonomous execution is not a future scenario. Enterprises are deploying systems today that act, transact, and commit resources without a person at the point of initiation. Each action carries the authority of the organization.

Existing control domains address pieces of the execution process, but none governs authority as a unified pre-execution function. Access controls, approvals, and audit remain useful within their domains. They do not resolve whether an autonomous system has sanctioned authority to perform a defined action in a defined relationship before execution occurs. The gap is not a weakness in any one control. The gap is the absence of authority infrastructure.

This paper defines that infrastructure. IGA separates identity, authorization, and enforcement into distinct functions built for autonomous environments. Governance decisions become infrastructure-level authority rules. The infrastructure resolves and enforces those rules before action occurs.

Execution governance now has an architecture. Enterprises deploying autonomous systems should ask whether their control environment includes one.

Appendix A: Implementation Considerations

This appendix sets out a practical implementation approach for bringing consequential automated execution within the IGA governance framework. It is intended to guide enterprise adoption and sequencing, not to prescribe a fixed technical specification.

Implementation should proceed through a governed enterprise process. That process begins with identifying where automated and autonomous execution occurs, establishing how consequential execution will be defined, determining which execution pathways fall within scope for IGA governance, and then implementing and maintaining the required control structure over time. The steps below provide a practical sequencing model for that work.

1. Inventory automated and autonomous execution pathways.

The organization should establish and maintain an enterprise inventory of systems, workflows, and processes in which execution may occur without case-by-case human intervention. The inventory should function as a living control record, not a one-time exercise.

2. Establish change-capture procedures.

Governance procedures should require review when new systems are introduced, workflows are redesigned, automation is expanded, or AI-enabled capabilities are materially modified. The purpose is to ensure that new or changed execution pathways are captured in the inventory when introduced, rather than identified after the fact.

3. Define consequential execution.

The organization should formally establish the criteria by which automated execution will be considered consequential. That determination should reflect the organization’s legal, financial, operational, regulatory, and stakeholder risk profile. Consequentiality may arise from monetary exposure, legal or regulatory effect, safety or security effect, records or reporting impact, contractual or relationship effect, customer or third-party consequence, or other forms of material organizational risk.

4. Determine which execution pathways fall within scope.

Each inventoried execution pathway should be evaluated against the organization’s defined criteria to determine its governance status under the IGA. The result should be a documented scope determination identifying which automated execution pathways are consequential and therefore subject to governed execution control.

5. Implement the IGA for in-scope pathways.

Organizations are not required to implement the IGA across all in-scope execution pathways at once. Implementation may begin with a limited set of consequential pathways and expand incrementally over time. For each pathway brought within scope, the organization should define the relevant role structure, authorization conditions, control dependencies, and execution pathway required to govern that action through the Identity Layer, Authorization Layer, and Enforcement Layer.

6. Maintain, review, and update over time.

The inventory, consequentiality criteria, scope determinations, and in-scope execution pathways should be reviewed and updated as systems, processes, and automation capabilities evolve. Implementation of the IGA is therefore not a one-time deployment exercise, but an ongoing governance program.

Implementation can proceed modularly, but the governance objective remains unified. The objective is to ensure that consequential automated execution is brought within a unified governance framework and remains subject to sustained enterprise oversight as systems and automation capabilities evolve.

Appendix B: Regulatory, Standards, Professional, and Legal Materials Reviewed

This appendix summarizes selected publicly available regulatory, supervisory, legislative, standards-based, professional, legal, and industry materials reviewed in connection with this paper as of the date of issuance. It is presented for contextual reference only and does not constitute legal, regulatory, or interpretive commentary.

Across these materials, recurring themes include formal accountability, documented decision rights, supervisory oversight, auditability, reproducibility, explicit rule governance, policy enforcement, and pre-execution control of consequential actions. Those themes correspond most directly to the IGA's structural separation of identity, authorization, and enforcement into distinct infrastructure-level control domains.

The materials reviewed informed the governance, accountability, oversight, control integrity, auditability, and execution-control themes reflected in the IGA set out in this paper. The references below are grouped by source category and issuing body. Bracketed references identify the applicable IGA control domain: I for Identity Layer, A for Authorization Layer, and E for Enforcement Layer.

U.S. Securities and Exchange Commission (SEC)

- SEC Roundtable on Artificial Intelligence in the Financial Industry (Mar. 27, 2025 / Apr. 9, 2025). AI governance in financial services. [A]
- Feedback on SEC Roundtable on AI (File No. 4-850) (Apr. 9, 2025). Public risk and governance concerns. [A]
- Remarks at the SEC AI Roundtable (Uyeda) (Mar. 27, 2025). AI risk oversight considerations. [A]
- Office Hours: Systemic Risk in Artificial Intelligence (Sep. 19, 2024). Systemic AI risk in markets. [A]
- Office Hours: Fraud and Deception in Artificial Intelligence (Oct. 10, 2024). Identity accountability and fraud prevention. [I, A]
- Artificial Intelligence and the Future of Investment Management (Feb. 3, 2026). AI integration in investment operations. [A]
- Conflicts of Interest Associated with Predictive Data Analytics (Federal Register) (Aug. 9, 2023). Conflict mitigation and rule constraints. [A]
- Remarks on AI Governance and Risk (FSOC Roundtable) (Mar. 4, 2026). Principles-based oversight, materiality, and continued human accountability in AI-enabled systems. [A]

Public Company Accounting Oversight Board (PCAOB)

- Staff Spotlight on Integration of GenAI in Audits (Jul. 3, 2024). Audit supervision and responsibility. [I, A]
- Staff Observations from Outreach on GenAI (Jul. 22, 2024). Accountability and audit controls. [A]
- Staff Priorities for 2025 Inspections (Dec. 8, 2024). Inspection focus areas. [A]
- Staff Report on 2025 Inspection Priorities (Dec. 9, 2024). Audit quality risk themes. [A]
- Staff Publications Index (ongoing). Repository of oversight publications. [A]

American Institute of CPAs (AICPA)

- Guidelines for Responsible Use of AI in Forensic & Valuation Services (Sep. 30, 2025). Professional accountability in AI usage. [I, A]
- Enhancing Audit Quality (EAQ) 2024 Highlights (Mar. 19, 2025). Audit quality in AI-enabled environments. [A]
- AI Auditing Strengthens Internal Controls (Feb. 10, 2026). Internal control implications of AI. [A]
- Artificial Intelligence Topic Hub (ongoing). Professional AI guidance repository. [A]

European Union

- Regulation (EU) 2024/1689 — Artificial Intelligence Act (Jun. 13, 2024 / Jul. 12, 2024). High-risk AI system obligations. [I, A, E]
- European Commission Proposal COM(2021) 206 (Apr. 21, 2021). Risk-based AI governance model. [A, E]

United States Government

- OMB Memorandum M-24-10 (Mar. 28, 2024). Federal AI governance requirements. [A]
- OMB Memorandum M-24-18 (AI Acquisition) (Sep. 24, 2024). AI procurement governance. [A]
- OMB Memorandum M-25-21 (Apr. 3, 2025). AI accountability and oversight. [A]
- Executive Order: Removing Barriers to American Leadership in Artificial Intelligence (Jan. 23, 2025). Federal AI policy direction. [A]
- America’s AI Action Plan (Jul. 10, 2025). National AI governance framework. [A]
- U.S. Treasury AI risk management resources for financial services (as reviewed as of the date of issuance). Cybersecurity, risk monitoring, and governance frameworks for AI deployment. [A]

Standards Bodies

- NIST AI Risk Management Framework (AI RMF 1.0) (Jan. 2023). AI risk management doctrine. [A]
- ISO/IEC 42001:2023 — AI Management Systems (Dec. 2023). AI management system requirements. [A]
- COSO, Achieving Effective Internal Control Over Generative AI (2026). Internal control considerations for generative AI systems. [I, A, E]

U.S. Legislative Monitoring

- Brennan Center, AI Legislation Tracker (ongoing). Federal AI legislative activity. [A]
- NCSL, Artificial Intelligence Legislation Summary (2025). State-level AI legislative themes. [A]

Professional, Legal, and Industry Materials

- Deloitte: Artificial Intelligence in the Boardroom: Five Considerations for Managing Risk in the AI Era. Board oversight, governance accountability, and control design. [I, A, E]
- Deloitte: AI and Risk Management. Enterprise risk and AI risk management frameworks. [A]
- PwC: Responsible AI Can Improve Business and Preserve Value. Supervisory controls and documentation for responsible AI. [I, A]
- EY: Responsible AI Principles. Accountability and assurance expectations for ethical deployment. [I, A]
- KPMG: Reimagining Controls Assurance with AI: Latest Insights for Leaders. AI assurance models and audit testing. [A]
- KPMG India: On the 2026 Board Agenda (Mar. 2026). Board oversight of AI strategy, governance guardrails, data governance, cybersecurity, and committee risk oversight as AI systems increase in autonomy and operational impact. [A]
- Bloomberg Industry Analysis: Big Four Firms Roll Out AI That Can Handle Routine Tasks Solo. Operational deployment risks as AI assumes routine professional tasks. [I, A]
- Thomson Reuters Legal: Implement AI Strategically — A Guide for Legal Departments. Policy enforcement and legal risk in AI-enabled legal operations. [A]
- LexisNexis Legal Insight: AI Risk Management in the United States: Looking Ahead. Compliance, defensibility, and documentation of AI decisions. [A]
- Skadden: The EU AI Act: What Businesses Need to Know. High-risk system obligations and regulatory analysis. [I, A, E]
- Accenture: AI and Your Operating Model: Radically New Ways of Working. AI-enabled operating model redesign and enterprise strategy. [A]
- ISDA: The Impact of Distributed Ledger Technology in Capital Markets. Distributed-ledger market infrastructure, tokenization, and governance, legal, and operational conditions for capital-markets adoption. [A, E]