THE ENTERPRISE AI GOVERNANCE FRAMEWORK

Govern AI at the point of execution

nuggetsLabs.com

CONTENTS

TABLE OF CONTENTS

01 - EXECUTIVE SUMMARY

EXECUTIVE SUMMARY

This document is for leaders responsible for deploying AI systems in production, teams evaluating AI vendors, and those who must ensure those systems remain controlled, auditable and accountable, including CISOs, CIOs, Chief Risk Officers, Heads of AI and Procurement.

AI systems are transitioning from generating content to executing actions across enterprise environments.

This shift introduces a new requirement, not only to secure AI systems, but to determine whether their actions are authorized to execute.

Traditional identity and access management governs who can access systems. That is a necessary control. It is not a sufficient one. IAM answers one question: can this actor reach this system? Action Governance answers a different question entirely: should this specific action happen right now? These are not the same question. For AI systems that analyze information, IAM is enough. For AI systems that execute actions, it is not.

Without Action Governance, organizations face a governance liability they cannot currently evidence.

If an AI system executes an unauthorized action, such as initiating a transaction, modifying infrastructure or accessing sensitive records, most enterprises today cannot prove what was authorized, by whom, or whether the action occurred within defined constraints. That is an unacceptable position as regulators, auditors and boards begin asking exactly these questions.

The Enterprise AI Governance Framework defines the capabilities required to ensure AI systems remain identifiable, authorized, constrained, verifiable and controllable.

The framework introduces three governance domains: Model Governance, System Governance and Action Governance, supported by Runtime Governance and Trust Infrastructure.

Action Governance is the new domain. It addresses the gap that existing frameworks do not: determining whether AI actions are authorized to execute, enforcing that decision in real time, and producing verifiable evidence of what occurred.

THE TRUST STACK

The framework is anchored by a trust stack that ensures AI actions are authorized and enforced across the full chain of responsibility:

Identity → Authority → Intent → Action

WHAT THIS ENABLES

Organizations that establish this control layer are positioned to deploy autonomous AI systems in regulated environments and demonstrate governance to regulators, auditors and boards on demand.

02 - ANALYST BRIEF

ANALYST BRIEF

> The stakes have changed.

In 2023, the primary concern with enterprise AI was whether models were accurate, biased or hallucinating. Those concerns remain valid, but they are no longer the primary governance risk.

The primary risk in 2026 is autonomous execution: AI agents and workflows, often without a human in the loop.

This is not a future scenario. It is the current state of enterprise AI deployment. By the end of 2026, the majority of new enterprise AI initiatives will involve systems with direct access to production infrastructure, live data and operational workflows, representing a significant shift from the analytical and generative deployments that dominated just two years prior. Yet the infrastructure to govern these systems has not kept pace. A recent report from the MIT Media Lab, *The GenAI Divide*, found that 95% of enterprise generative AI pilots are failing to deliver any measurable business return, what researchers call "pilot purgatory," driven in large part by the absence of trusted governance at the point of execution.

This shift is already exposing a consistent failure mode in enterprise AI systems. The financial exposure is real: IBM's 2024 Cost of a Data Breach Report places the average breach cost at $4.88M, and where AI systems are embedded in operational workflows, exposure compounds. Under the EU AI Act, non-compliance for high-risk AI systems carries fines of up to €30M or 6% of global annual turnover. Regulators are no longer treating AI governance as a future concern.

OBSERVED FAILURE MODE

Recent incidents and red-team research across major AI platforms have demonstrated that once AI systems are granted access to internal tools, data or infrastructure, they can be induced to perform actions outside their intended scope, including exposing sensitive data, executing unintended operations, or bypassing policy constraints through prompt manipulation or indirect input.

In these cases, the system does not fail at the level of access. It fails at the level of execution.

The AI system is operating within its granted permissions, but performing actions that were never explicitly authorized.

The result is a structural gap in existing governance models: organizations can control what systems can access, but not whether specific actions should be allowed to occur once that access is established.

The defining failure mode of agentic AI is not unauthorized access. It is unauthorized execution.

The control point has shifted from access to execution.

02 - ANALYST BRIEF (CONTINUED)

Enterprises must be able to determine, at the point of execution:

- Whether an AI action is authorized
- On whose behalf it is being performed
- What constraints govern whether it may execute
- Whether it should proceed, be constrained or be blocked

Traditional identity, security and compliance systems are not designed to make these decisions. They establish identity and access, but do not determine whether an action is authorized to execute across systems.

Consider this: if an AI agent in your environment executed an action that caused harm, such as a misconfigured network change, an unauthorized data transfer or an erroneous transaction, could you demonstrate what authority it was operating under?

Could you show a regulator exactly when the action occurred and whether it was within its authorized scope?

Could you produce a verifiable record of who delegated the authority for it?

For most enterprises, the honest answer is no.

The result is a governance gap between access control and action authorization.

The Enterprise AI Governance Framework defines the control model required to close this gap. It establishes how identity, delegated authority, intent and policy constraints combine to determine whether AI actions are authorized at the point of execution, and how those decisions are enforced and proven.

03 - THE CORE INSIGHT

THE CORE INSIGHT

AI governance has historically focused on models, ensuring they are safe, reliable and compliant. As AI systems begin executing actions across enterprise infrastructure, the control problem shifts.

The control point moves from who can access systems to whether actions are authorized to execute.

This is a fundamental change. In autonomous systems, identity and access are not sufficient. Governance must determine, in real time, whether an action should happen at all.

The clearest way to see this distinction: an AI agent may have legitimate access to a financial system, and still execute a transaction it was never authorized to make. Only Action Governance can stop it at execution.

This requires a new control model that combines identity, delegated authority, intent and policy constraints to determine whether AI actions are authorized at the point of execution.

TRADITIONAL ACCESS CONTROL ASKS

"Can this actor reach this system?"

ACTION GOVERNANCE ASKS

"Is this specific action, by this actor, under this authority, within these constraints, permitted to execute right now?"

04 - WHY IAM IS NOT SUFFICIENT FOR AUTONOMOUS AI

WHY IAM IS NOT SUFFICIENT FOR AUTONOMOUS AI

Identity and access management is essential infrastructure. Every enterprise needs it. But IAM was designed for a world where humans log in, perform actions, and log out. It governs the moment of access. It was never designed to govern the moment of execution.

As AI systems extend IAM into the agentic era, assigning identities to agents, issuing scoped tokens and enforcing delegated authority, they are solving the right problem for the wrong moment. Knowing who an agent is, and what systems it can reach, does not tell you whether a specific action it is about to take is authorized to proceed.

IAM governs access. Action Governance governs execution. These are different control points, and only one of them governs autonomous AI.

The distinction matters because autonomous AI agents do not just access systems. They act within them. An agent with legitimate, scoped, fully-credentialed access to a financial platform can still execute a transaction that was never authorized. IAM permitted it into the system. Only Action Governance can stop it at the moment of action.

The table below maps where each control layer operates and what it governs:

CONTROL QUESTION	TRADITIONAL IAM	ACTION GOVERNANCE
Core question	Can this actor access this system?	Should this specific action execute right now?
Control point	Authentication and session	Point of execution - before the action takes effect
Evaluates	Identity and access permissions	Identity, delegated authority, intent and constraints - together
Timing	Before session begins	At the moment of action - after access, before execution
Handles autonomous agents	Partially - identity only	Yes - identity, authority, intent and constraints
Produces verifiable evidence	Access logs	Cryptographic proof of authorized execution
Can block a specific action	No - governs access, not actions	Yes - allow, constrain or block at execution
Satisfies regulatory accountability	Partially	Yes - verifiable, auditable, on demand

> IAM and Action Governance are complementary, not competing. IAM establishes who the agent is. Action Governance determines whether the agent is permitted to act. Both are required. Neither is sufficient alone.

05 - THE AI GOVERNANCE GAP

THE AI GOVERNANCE GAP

Most organizations currently govern AI through model safety reviews, risk assessments and traditional IT security controls. These controls remain essential but were designed for AI systems that analyze information rather than act upon it.

Traditional identity and access management systems determine which actors can access enterprise systems. They do not determine whether a specific AI action is authorized to execute, under what constraints, or whether it should be blocked.

THE REGULATORY GAP

The gap is not theoretical. The EU AI Act, emerging financial services regulations and critical national infrastructure frameworks are all converging on a common requirement: AI systems that perform operational actions must be auditable, controllable and accountable at the point of execution. Governance frameworks that address only Model Governance and data access will not satisfy that requirement.

As AI systems begin performing actions across enterprise infrastructure, organizations must ensure those actions are authorized, constrained, enforceable, verifiable and controllable.

The Enterprise AI Governance Framework addresses this gap by introducing governance mechanisms that combine identity, delegated authority, intent and policy constraints to determine whether AI actions are authorized at the point of execution, and to enforce those decisions in real time.

06 - THE TRUST STACK

THE TRUST STACK

Safe deployment of autonomous AI systems requires a verifiable chain of responsibility for every action performed by an AI system. The framework defines this chain as a trust stack:

Identity -> Authority -> Intent -> Action

These primitives combine to determine whether an AI action is permitted to proceed, and whether it should be allowed, constrained or blocked.

01

IDENTITY

Who is acting

The verifiable identity of the actor performing an action. Actors may include humans, organizations, AI agents or machines. Identities must be cryptographically verifiable and linked to accountable principals.

02

AUTHORITY

On whose authority

Who authorizes the action and what permissions apply. Authority defines the rights delegated to an actor, including scope and limits. Authority must be enforceable at runtime and remain revocable.

03

INTENT

What is intended

The declared purpose and scope of a requested AI action. Intent provides the context required to determine whether an action is authorized to execute under delegated authority and policy constraints.

04

ACTION

What action is permitted to execute

The execution of an operational act and its controlled outcome. Actions must be authorized and enforced at the point of execution, with verifiable evidence that they occurred within defined constraints.

> Together, these primitives answer the four questions every governed AI action must satisfy: who is acting, on whose authority, what is intended, and what action is permitted to execute. If any of these primitives are missing, the action cannot be governed, only observed after the fact.

07, 08, 09 - GOVERNANCE DOMAINS

GOVERNANCE DOMAINS

The framework defines three governance domains that together ensure AI actions are authorized and controlled at the point of execution.

MODEL GOVERNANCE

Model Governance ensures AI models are safe, reliable and aligned with intended use. It is a pre-condition for safe AI deployment. It is not the control point.

SYSTEM GOVERNANCE

System Governance ensures AI systems interact safely with enterprise infrastructure. It establishes the boundaries of what AI systems can reach. It is not the control point.

ACTION GOVERNANCE

> Action Governance is to autonomous AI what identity and access management was to enterprise software. IAM established who could access systems. Action Governance establishes whether actions are permitted to occur. Every enterprise that deployed enterprise software eventually needed IAM. Every enterprise deploying autonomous AI will need Action Governance.

> Action Governance is the control point.

Action Governance determines whether AI actions are authorized to execute, under what authority and constraints, and whether they should proceed, be constrained or be blocked.

As AI systems move beyond analysis and begin executing tasks, organizations must ensure that every action is evaluated and enforced at the point of execution. Action Governance combines identity, delegated authority, intent and policy constraints to make that determination.

Action Governance is not a monitoring layer. It is an execution control point. Action Governance is evaluated as a policy decision before execution. It evaluates identity, delegated authority and intent. Only if all conditions are satisfied does the action proceed.

Key areas include AI actor identity, delegated authority, intent validation, policy enforcement at execution, constraint validation, real-time decisioning (allow, constrain, block), and intervention and revocation.

10 - RUNTIME GOVERNANCE

RUNTIME GOVERNANCE

Runtime Governance ensures that governance controls are enforced continuously while AI systems operate in production environments. It determines, in real time, whether AI actions are authorized and compliant.

Runtime Governance is the enforcement layer across all domains. It is not a monitoring capability. It is a decision-making capability. It enables actions to be permitted, constrained or blocked as they occur.

POLICY
Runtime policy enforcement at the point of action execution.

VALIDATE
Constraint validation before actions take effect.

DECIDE
Real-time decisioning: allow, constrain or block.

INTERVENE
Intervention and revocation of delegated authority.

MONITOR
Continuous monitoring of action execution.

11 - GOVERNANCE LIFECYCLE

GOVERNANCE LIFECYCLE

AI governance must operate across the full lifecycle of AI systems, with execution as the critical control point at every stage.

DISCOVER
Identify AI systems, agents and shadow AI usage across the organization.

EVALUATE
Assess models, systems and vendors, including risks associated with delegated authority and potential actions.

DEPLOY
Establish governance controls before production use, including identity, delegated authority, intent and policy constraints.

OPERATE
Determine and enforce, at the point of execution, whether AI actions are authorized, constrained or blocked.

EVOLVE
Continuously update policies, authority models and governance mechanisms as AI systems and their behaviors change.

AI GOVERNANCE LIFECYCLE

ENTERPRISE AI GOVERNANCE FRAMEWORK

A lifecycle model for governing AI systems in production

Most AI governance stops at models and access. This framework extends control to execution, ensuring every AI action is authorized, constrained and verifiable.

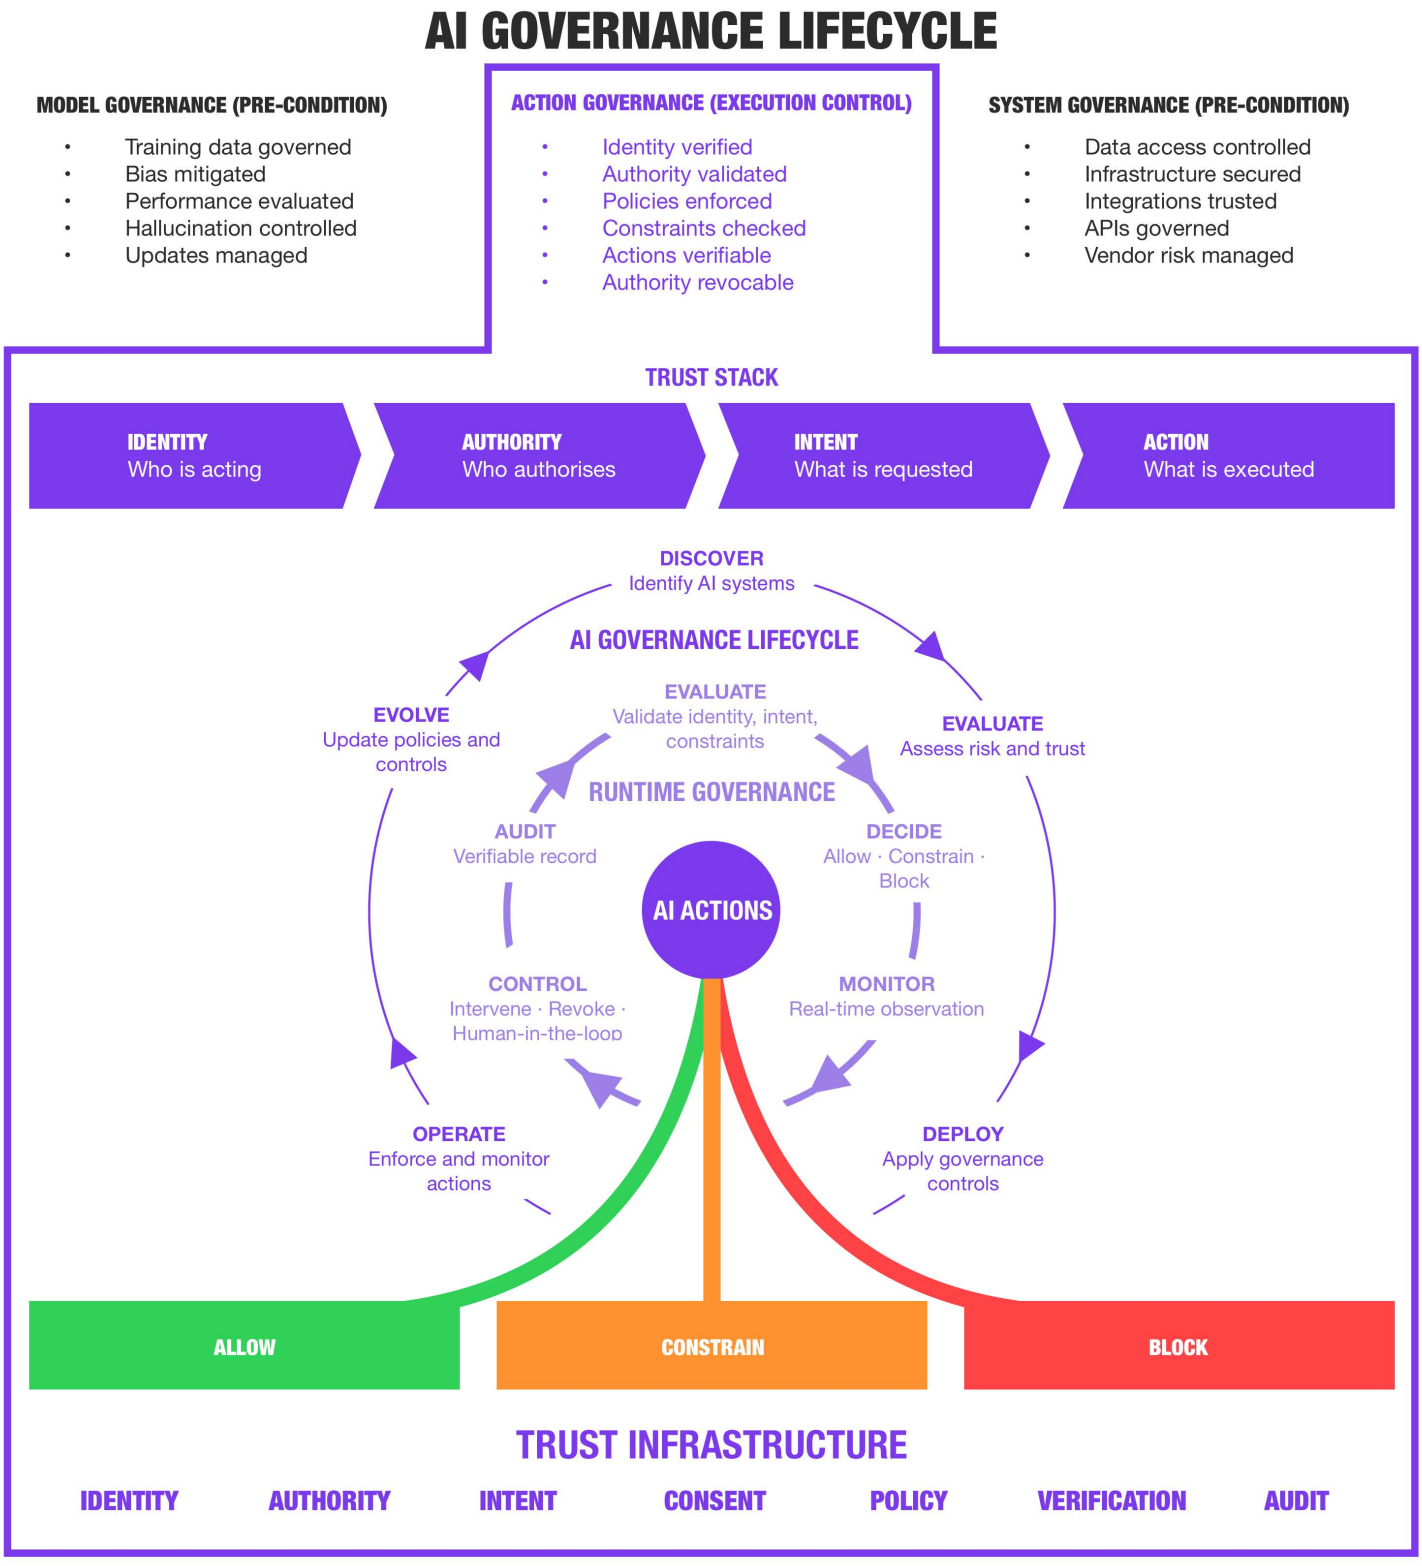

12 - AI RISK CLASSIFICATION

AI RISK CLASSIFICATION

Risk in autonomous AI systems is not primarily a function of model capability. It is a function of what those systems are allowed to do, and what happens if those actions are incorrect or uncontrolled.

Risk increases as AI systems move from generating outputs to executing actions across enterprise systems. The framework classifies AI systems based on the authority they are granted and the potential impact of their actions.

- **Low Risk - Advisory Systems**
  AI systems that generate insights or recommendations but do not execute actions. Examples include summarization tools, research assistants and decision-support systems. Primary risks relate to accuracy and bias, with limited direct operational impact.

- **Operational Risk - Workflow Influence**
  AI systems that influence operational processes but do not directly execute actions. Examples include systems that recommend approvals or generate workflow instructions. Execution typically remains under human control.

- **High Risk - System Execution**
  AI systems capable of executing actions within enterprise systems, modifying data, triggering workflows or interacting with operational platforms. These systems require governance mechanisms that determine, at the point of execution, whether actions are authorized, constrained or blocked.

- **Critical Risk - Autonomous Execution**
  AI systems capable of initiating transactions, controlling infrastructure or operating without human intervention. These systems require continuous Runtime Governance to ensure every action is authorized, policy-constrained and enforceable at the point of execution, with the ability to intervene or revoke actions in real time.

Most enterprise AI deployments in 2026 are entering the High Risk and Critical Risk tiers without the governance infrastructure those tiers require. That is the gap this framework addresses.

13 - ENTERPRISE AI GOVERNANCE ROLES

ENTERPRISE AI GOVERNANCE ROLES

Effective governance of autonomous AI systems requires clearly defined roles responsible for defining, authorizing and enforcing AI actions at the point of execution.

AI System Owner
Responsible for the overall operation of the AI system, including how it interacts with enterprise infrastructure and executes actions.

Model Owner
Responsible for model performance, evaluation and updates.

Policy Authority
Defines the rules and constraints governing what AI systems are permitted to do, including acceptable actions, conditions and boundaries.

Security Oversight
Ensures AI systems meet enterprise security standards, including identity, access and infrastructure controls.

Compliance Oversight
Ensures AI systems operate in accordance with regulatory, legal and organizational requirements.

Runtime Governance Authority
This role does not exist in most organizations today. It is responsible for determining and enforcing, at the point of execution, whether AI actions are authorized, constrained or blocked. As AI systems become more autonomous, this becomes one of the most consequential governance roles in the enterprise.

14 - AI GOVERNANCE INFRASTRUCTURE PRIMITIVES

AI GOVERNANCE INFRASTRUCTURE PRIMITIVES

Enterprise AI governance depends on foundational infrastructure primitives that combine to determine whether AI actions are authorized to execute, and how those decisions are enforced at runtime.

Identity
Cryptographically verifiable identity for all actors, including humans, organizations, AI agents and machines.

Authority
Delegated rights defining what actions an entity is permitted to perform, including scope, limits and revocation.

Intent
Explicit declaration of the action requested, including context and expected outcome.

Consent
Approval artifact defining what actions are permitted and on whose behalf they may be performed.

Policy
Rules and constraints governing how and when actions may be executed, enforced consistently across systems.

Enforcement
Runtime capability to determine whether actions are authorized and to allow, constrain or block execution.

Verification
Cryptographic proof that actions were authorized and executed within defined constraints.

Audit
Records of decisions, actions and outcomes to support accountability, investigation and compliance.

15 - PRIVACY-PRESERVING TRUST INFRASTRUCTURE

PRIVACY-PRESERVING TRUST INFRASTRUCTURE

Modern AI governance requires trust infrastructure capable of enforcing identity and execution, without creating new privacy liabilities in the process.

Rather than relying on centralized logging or traditional identity systems, these architectures use cryptographic and verifiable mechanisms to enable actions to be authorized and enforced without exposing unnecessary data. Examples include:

- Verifiable credentials
- Selective disclosure proofs
- Cryptographic delegation chains
- Privacy-preserving audit evidence
- Consent-bound authorization artifacts

16 - ALIGNMENT WITH INDUSTRY STANDARDS

ALIGNMENT WITH INDUSTRY STANDARDS

The Enterprise AI Governance Framework complements existing AI governance and regulatory frameworks. It defines Action Governance, the execution control layer that existing frameworks do not.

NIST AI RISK MANAGEMENT FRAMEWORK

The framework maps to all four NIST AI RMF core functions. It extends NIST by introducing Action Governance, the execution control layer NIST's Manage function does not address: determining whether individual AI actions are authorized to execute in real time.

Govern	Governance lifecycle and organizational roles
Map	Discovery and evaluation of AI systems
Measure	Model Governance and risk assessment
Manage	System Governance and Runtime Governance

EU AI ACT

The framework supports the EU AI Act's core requirements. It addresses the compliance challenge the Act creates but does not solve: how to demonstrate, on demand, that a high-risk AI system's actions were authorized, constrained and within defined parameters at the time of execution.

Risk classification	AI risk tiers based on operational authority and action impact
Transparency	Model Governance and explainability
Monitoring	Runtime Governance and continuous oversight
Accountability	Action Governance and verifiable execution

ISO AI GOVERNANCE STANDARDS

The framework aligns with ISO AI governance standards across all four domains. It extends ISO guidance by introducing the execution control layer, determining whether AI actions are permitted to execute, rather than assessing risk before or after the fact.

AI management	Governance lifecycle and organizational structure
Risk management	Model Governance and evaluation
Operational oversight	System Governance and Runtime Governance
Accountability	Action Governance and traceability

17 - CONCLUSION

CONCLUSION

Most existing AI governance frameworks focus on model development, safety and risk management. As AI systems begin executing actions across enterprise infrastructure, governance must extend beyond models to determine whether those actions are authorized to execute.

The Enterprise AI Governance Framework extends existing governance models by ensuring that AI systems remain identifiable, authorized, constrained, enforceable, verifiable and controllable.

Implementing these capabilities requires infrastructure that combines identity, delegated authority, intent and policy constraints to determine whether AI actions are authorized at the point of execution, and to enforce those decisions consistently across systems.

The question facing every enterprise deploying autonomous AI is not whether they need this control layer. They do. The question is whether they establish it before something goes wrong, or after.

Organizations that govern AI at the point of execution will be able to deploy autonomous systems at scale, satisfy regulatory scrutiny, and demonstrate accountability to boards and auditors on demand. Those that do not will face the same question every time an AI system acts without provable authorization: how do you know it was allowed to do that?

Action Governance is the answer. The Enterprise AI Governance Framework provides the structure to build it: a governance lifecycle, a risk classification model, defined roles, infrastructure primitives, and a trust stack that makes every AI action auditable and accountable.

Organizations can begin by assessing where their AI systems sit on the risk classification scale, identifying the governance gaps in their current infrastructure, and establishing the identity, authority and policy controls required before autonomous AI systems execute consequential actions.

Where to start: three steps

1. Classify.
Map your active AI deployments against the risk tiers in this framework. Identify which systems are already operating at the High or Critical tier, systems that execute actions, not just generate outputs.

2. Gap-assess.
For each High or Critical system, ask three questions: can you verify the identity of the AI actor, prove the authority it was operating under, and produce tamper-resistant audit evidence on demand? If the answer to any is no, that is your governance gap.

3. Prioritize.
Establish identity and delegated authority controls for your highest-risk systems first, before expanding to full policy enforcement and runtime governance. IAM is the foundation; Action Governance is the next layer. Both are required. Start where the exposure is greatest.

18 - PROCUREMENT QUESTIONS

PROCUREMENT QUESTIONS FOR EVALUATING AI SYSTEMS

Enterprise teams evaluating AI systems should ensure vendors can demonstrate governance capabilities across identity, delegated authority and execution control. These questions assess whether AI systems can safely operate in environments where they perform actions, not just generate outputs.

IDENTITY AND ACCOUNTABILITY

1. How are AI actors identified within the system?
2. Can AI agents be linked to accountable human or organizational principals?
3. Are identities cryptographically verifiable across systems and services?

AUTHORITY AND DELEGATION

4. How is authority delegated to AI systems or agents?
5. Can delegated authority be restricted, revoked or time-bound?
6. Can the system evaluate authority at the point of execution before actions occur?

POLICY ENFORCEMENT AND CONSTRAINTS

7. How are policies enforced when AI systems perform actions?
8. Can actions be evaluated against policy constraints at the point of execution?
9. Can the system allow, constrain or block actions in real time?

EXECUTION CONTROL

10. How does the system determine whether an AI action is authorized before it executes?
11. Are identity, authority, intent and policy evaluated together at runtime?
12. Can unauthorized or out-of-scope actions be prevented before execution?

VERIFICATION AND EVIDENCE

13. What evidence exists that AI actions were authorized and executed within defined constraints?
14. Can organizations independently verify both the authorization decision and execution outcome?
15. Are decisions and actions recorded in a verifiable and tamper-resistant manner?

RUNTIME GOVERNANCE

16. Do governance controls remain active while AI systems operate autonomously?
17. Can organizations intervene, constrain or revoke actions in real time?
18. Are governance controls consistently enforced across systems, APIs and cloud environments?