WHITE PAPER

Applying IAMM to Artificial Intelligence

A maturity-based approach to governing AI risk, assurance, and trust

An executive guide to applying the Integrated Assurance Maturity Model (IAMM) across the AI lifecycle.

AI adoption is moving faster than many organizations can govern it. IAMM provides a way to see where assurance is fragmented, where it is aligned, and where it has become part of how the enterprise operates.

Patrick M. Hayes

Integrated Assurance®

Executive white paper • September 2026

EXECUTIVE SUMMARY

AI assurance has to mature with AI adoption

Artificial intelligence is becoming embedded in customer engagement, financial systems, logistics, infrastructure, decision-making, and automation. The same qualities that make AI useful also create new forms of exposure. Models can be opaque, data pipelines can be fragile, autonomous systems can exceed intended boundaries, and adversaries can exploit weaknesses faster than periodic controls can adapt.

The consequence is larger than technical disruption. A compromised chatbot, manipulated dataset, deepfake campaign, or uncontrolled AI agent can undermine trust, create regulatory exposure, and leave leaders accountable for outcomes they cannot adequately explain.

The Integrated Assurance Maturity Model (IAMM) provides a structured way to address that problem. It does not replace established standards. It connects governance, architecture and engineering, IT and operations, risk and compliance, metrics and reporting, and culture and collaboration so that assurance can operate across the full AI lifecycle.

Applied to AI, IAMM helps leaders answer three practical questions:

- **Where are we now?** Understand maturity across the six assurance domains rather than relying on isolated controls or policies.
- **Where are the gaps?** Identify where AI risk is outpacing governance, engineering, operational visibility, compliance, measurement, or cultural accountability.
- **What should improve next?** Use maturity as a roadmap for deliberate progress toward embedded and institutionalized assurance.

> **The goal is not to make AI risk-free. It is to create enough visibility, evidence, accountability, and operational discipline to know when AI can be trusted and when intervention is required.**

Traditional controls were built for a more predictable environment

AI does not behave like the information systems most governance models were designed around. Its logic can be probabilistic, its behavior depends heavily on data, and its operating state can change as models, inputs, integrations, and real-world conditions evolve.

A model can perform well by conventional technical measures and still create problems involving explainability, fairness, accountability, resilience, or regulatory defensibility. That is why trust cannot be inferred from accuracy alone.

Three conditions make AI assurance different

- **Opacity.** Organizations may be responsible for decisions whose underlying reasoning is difficult to explain or reproduce.
- **Data fragility.** Poor, biased, manipulated, or inadequately governed data can influence thousands of downstream decisions.
- **Autonomy.** AI systems and agents can adapt, act, escalate, and interact with tools in ways that exceed static control assumptions.

Periodic audit and siloed oversight are not enough for systems that change continuously. AI assurance must connect governance to engineering, compliance to operations, and culture to accountability.

AI creates exposure across data, interfaces, models, people, and autonomous action

AI RISK	EXPOSURE	IAMM RESPONSE
Model opacity	Outputs may be difficult to explain or justify.	Explainability expectations, human oversight, transparency measures
Data fragility	Poisoned, biased, incomplete, or manipulated data undermines model integrity.	Provenance, validation, lineage, bias and drift monitoring
Autonomous behavior	Agents may adapt or escalate beyond intended boundaries.	Escalation protocols, review checkpoints, adversarial simulation
Prompt injection	Natural-language inputs can manipulate model behavior or connected workflows.	Adversarial testing, interface hardening, behavioral monitoring
Synthetic media	Deepfakes and voice cloning undermine identity and reputation.	Detection, verification, rehearsed response, communications coordination
Shadow AI	Unsanctioned tools create data leakage and accountability gaps.	AI inventories, sanctioned alternatives, monitoring, acceptable-use practices
Regulatory complexity	AI obligations span jurisdictions, sectors, and evolving standards.	Control mapping, lifecycle documentation, continuous readiness

The AI perimeter is no longer a firewall or application boundary. It is a mesh of data flows, interfaces, models, external services, autonomous actions, and human decisions.

Six domains create one view of assurance

IAMM treats AI assurance as an enterprise discipline. Each domain addresses a different part of the operating model, but the value comes from understanding how those domains work together.

Governance

Make AI accountability visible and enforceable. Establish ownership, oversight, model governance, human review expectations, and decision authority.

Architecture & Engineering

Build assurance into model design, training pipelines, deployment architecture, data provenance, reproducibility, validation, and interface resilience.

IT & Operations

Maintain visibility into deployed AI, shadow AI, telemetry, anomalies, approved platforms, inventories, and operational behavior over time.

Risk, Compliance & Audit

Connect AI adoption to enterprise risk, legal obligations, auditability, documentation, regulatory readiness, and evidence of control performance.

Metrics & Reporting

Measure more than accuracy and throughput. Track indicators that reveal drift, explainability, resilience, fairness, control integrity, and trust.

Culture & Collaboration

Create shared accountability across business, engineering, security, legal, risk, and compliance. Make responsible AI behavior part of normal work.

> **No single domain can compensate indefinitely for weakness in another. Strong engineering with weak governance, or strong policy with weak operational visibility, still leaves the organization exposed.**

From fragmented oversight to institutionalized assurance

IAMM uses five maturity levels to show how assurance becomes progressively more coordinated and embedded. The levels are guideposts, not a claim that every domain should advance at the same speed.

Level 1 - Fragmented

AI adoption occurs in silos. Shadow AI is common. Data provenance, oversight, escalation, and assurance practices are inconsistent or absent. Trust is largely assumed.

Level 2 - Defined

Policies, awareness, basic logging, and initial review practices begin to appear, but implementation remains inconsistent and governance is still largely reactive.

Level 3 - Aligned

Cross-functional oversight takes shape. AI systems are inventoried, critical models receive defined review, provenance and drift monitoring improve, and trust-related metrics begin to enter governance discussions.

Level 4 - Embedded

Assurance is integrated into normal workflows. Red teaming, interface testing, automated validation, shadow AI detection, continuous monitoring, and auditable human intervention become routine.

Level 5 - Institutionalized

AI assurance becomes an organizational capability. Governance is dynamic, third parties are included, compliance is continuous, board reporting includes trust and resilience, and accountability shapes strategic decisions.

> **Maturity is rarely uniform. IAMM is most useful when it exposes divergence between domains instead of hiding it behind a single enterprise score.**

The gaps between domains are often where risk accumulates

An organization can have strong board oversight and weak engineering evidence. It can have sophisticated model validation while employees continue to use unapproved AI tools. It can maintain an AI inventory but lack metrics that tell executives whether assurance is actually working.

These mismatches matter because AI risk crosses organizational boundaries. A weakness in culture can bypass an engineering control. A weakness in metrics can leave governance blind. A weakness in operations can allow a sanctioned model to become an unmanaged dependency after deployment.

Use the maturity profile to expose divergence

- Assess each domain independently rather than forcing one overall maturity label.
- Identify the lowest-maturity capabilities that materially affect high-consequence AI use.
- Look for dependencies between strong and weak domains.
- Direct investment toward gaps that reduce the credibility of otherwise mature controls.
- Reassess as AI adoption, vendors, models, regulations, and business use change.

The objective is not perfect symmetry. It is conscious progress with clear visibility into where assurance remains weaker than the business assumes.

IAMM connects existing frameworks to operational AI assurance

IAMM is not intended to displace the standards organizations already use. Its role is connective. It helps translate governance principles, management-system requirements, enterprise risk practices, and adversary knowledge into a coordinated operating model for AI.

FOUNDATION	ROLE WITH IAMM
NIST Cybersecurity Framework 2.0	Provides governance and security functions. IAMM extends those foundations into continuous AI-specific risk monitoring and operational assurance.
ISO/IEC 42001	Defines requirements for AI management systems. IAMM helps embed those expectations into engineering, operations, provenance, validation, and human oversight.
COBIT and COSO ERM	Provide governance and enterprise risk principles. IAMM connects those principles to technical assurance and cross-functional accountability.
MITRE ATT&CK;	Describes adversary tactics and techniques. IAMM uses that threat understanding to shape organizational countermeasures across technology, governance, data, and culture.

The practical value is not another control catalog. It is a way to make existing expectations work together across the lifecycle of AI.

Move from visibility to institutionalized trust

The source work describes a deliberate adoption path. It begins by understanding the current state, then prioritizes the gaps that matter most before embedding, validating, and scaling assurance.

1. Assess

Baseline maturity across all six IAMM domains. Identify AI systems in use, sanctioned and unsanctioned activity, data provenance, human oversight, metrics, and accountability.

2. Identify

Prioritize the highest-exposure gaps. Shadow AI, weak data integrity, high-consequence autonomous workflows, or regulatory readiness may require attention first.

3. Operationalize

Translate policy into working controls. Embed validation, explainability expectations, monitoring, escalation, human review, and operational evidence into workflows.

4. Integrate

Connect AI assurance to enterprise governance, security, risk, compliance, engineering, operations, and reporting so it does not remain a specialist initiative.

5. Validate

Test whether controls work. Use red teaming, prompt injection and model-inversion scenarios, data-quality stress tests, compliance drills, and exercises of human escalation.

6. Scale

Institutionalize assurance across the enterprise and ecosystem. Extend expectations to vendors and partners, include trust measures in board reporting, and sustain regulatory readiness.

> **The roadmap is not a rigid checklist. Organizations can move at different speeds, but progress should remain visible, deliberate, and connected to business consequence.**

What maturity changes in practice

Prompt injection in a customer chatbot

A customer-facing AI is manipulated into revealing sensitive operational details. At lower maturity, the event is treated as an isolated technical flaw. As maturity improves, red teaming, behavioral analytics, interface protections, threat intelligence, governance oversight, and board visibility turn the same risk into a managed assurance problem.

Shadow AI in financial services

Employees use consumer AI tools with confidential information. Early maturity leaves the behavior largely invisible. More mature organizations combine awareness with sanctioned alternatives, detection, central inventories, monitoring, and enterprise-wide governance so productivity does not require surrendering accountability.

Deepfake fraud

Synthetic voice or video is used to impersonate leadership and influence financial or public decisions. Mature assurance replaces assumed authenticity with verification, detection, rehearsed response, legal and communications coordination, and explicit responsibility for trust.

Agentic AI and adaptive attack activity

Autonomous adversaries can sustain reconnaissance, adapt attack paths, and exploit interfaces with less human intervention. Higher maturity requires continuous telemetry, adaptive defense, red teaming, clear escalation, and governance that treats autonomous behavior as an enterprise risk rather than only a technical event.

> **Across the scenarios, the technology changes. The assurance problem does not: organizations need visibility, evidence, ownership, coordinated response, and the ability to intervene.**

AI governance becomes credible when leaders can see assurance operating

Executives and boards do not need to manage model internals, but they do need a reliable view of where AI is being used, which uses can create material consequences, who owns the risk, what evidence supports trust, and how quickly the organization can intervene when conditions change.

Questions leadership should be able to answer

- Do we know where AI is being used across the organization and by critical third parties?
- Which AI uses can materially affect customers, finances, operations, regulatory obligations, reputation, or safety?
- Where is human review required, and is that intervention auditable?
- Can we trace the data, model version, inputs, outputs, approvals, and downstream actions behind a consequential decision?
- Which IAMM domains are weaker than the others, and what business exposure does that create?
- Are trust, resilience, and control performance visible alongside AI performance and cost?
- What would cause us to restrict, retrain, replace, or retire an AI capability?

A mature AI program is not defined by how much AI the enterprise deploys. It is defined by whether the organization understands the trust it is placing in those systems and can defend that trust with evidence.

CONCLUSION

Build AI on a foundation of trust

AI is becoming an enterprise capability rather than a discrete technology project. That changes the governance problem. Assurance cannot remain a review performed after deployment or a collection of policies that operate separately from engineering and operations.

IAMM provides a way to make that transition visible. Its six domains show where assurance must exist. Its five maturity levels show how fragmented practices can become aligned, embedded, and ultimately institutionalized. Its value is greatest when leaders use the maturity profile to expose gaps between organizational confidence and operational reality.

The objective is not uniform Level 5 maturity everywhere. It is to understand the level of assurance appropriate to the consequence, maintain visibility into unevenness, and improve deliberately as AI becomes more important to the business.

> **Trust in AI cannot be assumed from performance. It has to be built through governance, engineered into systems, monitored in operations, tested against risk, made visible through metrics, and reinforced through culture.**

About the source work

This white paper is a condensed executive treatment of Patrick M. Hayes's work on applying the Integrated Assurance Maturity Model to artificial intelligence. The full work examines AI as a trust challenge, the expanding threat surface, IAMM for AI, maturity levels, practical adoption, case scenarios, strategic implications, and building AI on a foundation of trust.

About Patrick M. Hayes

Patrick M. Hayes is an author, strategist, executive advisor, and keynote speaker whose work focuses on Business Survivability, Integrated Assurance, cybersecurity, AI, governance, and organizational risk. His work examines how technology, operations, trust, and decision-making interact under conditions of rapid change and uncertainty.

patrickmhayes.com | info@integratedassurance.co