

THE AI EXECUTION GAP

Why 94% of Enterprises Fail to Scale AI

Enterprise AI Production Readiness: The 2025 Execution Playbook


A diagnostic and capital allocation framework for enterprise AI success.

A Strategic Assessment of Enterprise AI Capabilities | 2024–2025 Edition

ENTERPRISE AI STRATEGY MAP


The AI Execution Gap — 2025 Edition

EARS FOUNDATION — Assess your maturity at each layer before committing capital to any capability.



CAPABILITY INVESTMENT BUCKETS — Classified by production readiness and time-to-value.

DEPLOY NOW <i>Proven ROI • Clear path to production • <6 months TTV</i>	Software Eng AI (Score 92)	Predictive Maint. (Score 90)	Fraud Detection (Score 88)	Generative AI (Score 85)	DSLMS (Score 80)
STRATEGIC BUILD <i>6–24 month horizon • Build EARS foundation first</i>	AI TRISM (Score 75)	Supply Chain AI (Score 70)	Agentic AI (Score 65)		
DEFER <i>24M+ horizon • Monitor open-weight model trajectory</i>	Multimodal AI (Score 45)	Multiagent Systems (Score 30)			



THE GENAI DIVIDE — Only ~5–10% of enterprises (varies by study) are capturing meaningful EBIT impact from AI. The gap is structural.

PACESETTER (~5–10%)

Significant EBIT impact • Proprietary data • EARS maturity

THE GAP (~90–95%)

POC Trap • Data Swamp • Governance Vacuum

"AI success is not a technology problem. It is a capital allocation problem under uncertainty."

Executive Summary

Enterprise AI has entered its execution phase. The models work. The organizations don't.

Enterprise AI has entered its execution phase. 72–78% of organizations report AI deployment in at least one function, yet only ~5–10% are achieving meaningful EBIT impact (varies by study and definition). Total worldwide AI spending is projected to approach ~\$1–1.5 trillion by the end of 2025, depending on definition of AI spend (IDC). The capital is flowing. The ROI is not for most.

3 Hard Truths for CXOs

- POCs don't fail. Scaling does. The 30–42% failure rate for AI initiatives concentrates in the transition from pilot to production, driven by data quality gaps, governance absence, and organizational inertia.
- Most AI systems fail silently before they fail visibly. Models degrade without triggering alerts, causing error rates to increase significantly (often 20–40% in poorly monitored systems) if left uncalibrated beyond six months.
- AI is not a technology transformation. It is a capital allocation problem under uncertainty. Organizations investing 70% of AI resources in people and processes — not just technology — are 4x more likely to reach the Pacesetter stage.

3 Decisions CEOs Must Make Now

- Prioritize Vertical AI: Precision-focused models integrated with proprietary data deliver up to 4x higher ROI than generic horizontal tools.
- Mandate Human-in-the-Loop (HITL): As AI handles repetitive execution, the burden of strategic oversight on human talent intensifies, not diminishes.
- Govern before you scale: Deploy the AI TRISM framework before expanding mission-critical workloads. The regulatory window is narrowing.

How Readiness Scores Are Calculated

Every capability in this report carries a Readiness Score (0–100). Scores are derived from four weighted dimensions:



Technical Maturity assesses production deployment evidence and toolchain stability. ROI Clarity reflects the quality and consistency of published financial outcomes. Adoption Level captures current enterprise deployment breadth. Risk Profile accounts for failure rate, adversarial risk, and regulatory exposure.

Strategic Capability Grouping

● IMMEDIATE ROI — Deploy Now

- Software Engineering AI (Score: 92) • Predictive Maintenance (Score: 90)
- Fraud Detection / Risk (Score: 88) • Generative AI / RAG (Score: 85)
- Domain-Specific LMs / DSLMs (Score: 80)

● STRATEGIC BUILD — Plan for 6–24 Months

- AI TRiSM (Score: 75) • Supply Chain AI (Score: 70) • Agentic AI (Score: 65)

● EXPERIMENTAL — Defer 24M+

- Multimodal AI (Score: 45) • Multiagent Systems (Score: 30)

The Enterprise AI Readiness Stack (EARS)

All 10 capabilities map to five foundational layers that determine whether an AI initiative succeeds or fails at enterprise scale. Capital decisions should start here, not at the model layer.

Layer 1: DATA — The Foundation

Quality, accessibility, and governance of training and inference data.

Key risk: Data readiness is the single most cited root cause of AI project failure (Gartner, 2025). No model selection decision compensates for a weak data layer.

Layer 2: MODELS — The Engine

Selection between general-purpose LLMs and domain-specific fine-tuned models.

Key risk: General models underperform in high-precision environments, achieving 65–75% accuracy vs. 95%+ for DSLMs on niche tasks.

Layer 3: SYSTEMS — The Infrastructure

Integration with legacy systems, APIs, and real-time data pipelines.

Key risk: 68% of CTOs cite legacy fragmentation as the top adoption blocker. Systems debt compounds model performance issues.

Layer 4: OPERATIONS — The Execution Layer

Monitoring, retraining cadence, incident response, and ROI measurement frameworks.

Key risk: Silent model degradation is the most common post-launch failure mode. Operations maturity is the primary differentiator between Pacesetters and laggards — not model selection.

Layer 5: GOVERNANCE — The Control Plane

Policy enforcement, regulatory compliance, explainability, and audit trails.

Key risk: Only 23% of IT leaders are confident in their AI governance capability. Gartner predicts a 30% increase in AI-related legal disputes by 2028.

The 5 AI Failure Archetypes

Failure analysis across enterprise AI deployments reveals five recurring patterns. Identifying your dominant archetype is the highest-leverage diagnostic available to a technology leader.

Archetype 1: The POC Trap

The most common failure mode. Organizations build impressive demos that lack the data pipelines, governance structures, and retraining cadence needed for production. Gartner predicts 30% of GenAI projects will be abandoned after POC by the end of 2025. Symptom: endless pilot loops with no production deployment date.

Archetype 2: The Data Swamp

Initiatives stall because data assets are fragmented, poorly labelled, or inaccessible across siloed systems. This is the most common and most preventable failure mode. Symptom: high compute spend, low model accuracy, repeated data remediation cycles with no progress.

Archetype 3: Cost Explosion

Inference costs for general-purpose LLMs (\$0.02–\$0.12 per token, rapidly declining as of 2025–2026) create unsustainable operational overhead at scale. Multi-agent architectures consume 1.6x–6.2x more tokens than single-agent workflows, prompting finance teams to kill projects mid-rollout. Symptom: ROI deteriorates sharply between months 6 and 18.

Archetype 4: Silent Degradation

Models erode in accuracy due to data drift and distribution shifts, often without triggering standard monitoring alerts. Error rates compound significantly in models left uncalibrated beyond six months — and failures typically surface first through customer complaints, not internal dashboards. Symptom: quality complaints precede internal detection by weeks.

Archetype 5: The Governance Vacuum

Organizations deploy AI at scale before establishing the risk controls, audit trails, and regulatory compliance frameworks required for sustained operation. Gartner predicts a 30% increase in AI-related legal disputes by 2028. Symptom: rapid capability deployment followed by forced rollbacks or regulatory remediation.

3 Contrarian Insights

The following perspectives diverge from prevailing industry consensus. They are grounded in deployment data and intended to stress-test assumptions that most enterprise AI roadmaps currently treat as settled.

CONTRARIAN VIEW

Most enterprises should NOT build custom LLMs. The economics rarely justify it outside of the top 50 global technology companies. For 95% of enterprises, fine-tuning an open-source base model (7B–13B parameters) on proprietary data delivers superior ROI, faster time-to-value, and lower operational risk than building from scratch. Custom LLM projects consuming \$10M+ budgets are frequently a status decision, not a strategic one.

CONTRARIAN VIEW

Agentic AI is overfunded relative to its maturity. The gap between investor enthusiasm and production reality is wider for agentic systems than for any other capability in this report. Early-stage deployments report high failure rates (often cited in the 60–90% range depending on scope), combined with \$500K+ annual orchestration overhead for poorly scoped deployments, suggests that current capital allocation to agentic AI in most enterprises is ahead of the organizational capability required to make it work. The right investment now is in EARS Layer 3–4 maturity, not agent orchestration.

CONTRARIAN VIEW

TRiSM will slow down AI adoption before it accelerates it. In the near term (2025–2027), the operationalization of AI governance frameworks will act as a brake on deployment velocity as organizations inventory models, establish audit trails, and respond to regulatory pressure. Enterprises treating TRiSM as a growth enabler today will experience this slowdown as a controlled deceleration. Those who delay governance will face an uncontrolled one.

The AI ROI Equation

The following formula distills the core dynamics of enterprise AI investment into a single, memorable framework. It is not a spreadsheet model — it is a strategic diagnostic. Each variable is a lever CXOs can act on.

$$\text{AI ROI} = (\text{Data Quality} \times \text{Operational Maturity} \times \text{Adoption Depth}) \div (\text{Model Cost} \times \text{Risk Exposure})$$

How to Use This Equation

- **Data Quality (numerator):** The single highest-leverage variable. A 10% improvement in data quality typically yields a 20–30% ROI improvement across all deployed capabilities — consistently the highest-return investment available to a technology leader.
- **Operational Maturity (numerator):** Monitoring, retraining cadence, and incident response. Organizations in the top quartile of operational maturity achieve 4x higher sustained ROI than those with strong launch-day performance but weak ongoing ops.
- **Adoption Depth (numerator):** Shallow adoption (5–10% of target users) produces shallow ROI. Real impact emerges at ≥60% active adoption with workflow integration. This is the most commonly underinvested variable.
- **Model Cost (denominator):** The variable most actively managed by finance teams. General-purpose LLM inference at \$0.02–\$0.12/token (rapidly declining as of 2025–2026) becomes the dominant cost at scale. DSLMs reduce this denominator by 40–60%.
- **Risk Exposure (denominator):** Includes adversarial risk, regulatory risk, and silent degradation risk. Unmanaged risk compounds over time — a TRISM-mature organization effectively reduces this denominator, multiplying ROI without changing any other variable.

The fastest path to higher AI ROI is usually not a better model. It is a better numerator.

Deep Case Study: Inside a Failed AI Transformation

The following composite case study is drawn from documented patterns across multiple enterprise AI transformations in the financial services and insurance sectors. It maps the failure arc to the EARS framework and the 5 Failure Archetypes, and shows the remediation path.

Case Profile	
Organization type:	Mid-size European insurance group (~8,000 employees, €4.2B GWP)
AI investment:	€12M over 24 months across 4 concurrent initiatives
Announced ambition:	"AI-first claims processing by Q4 Year 2."
Actual outcome:	All 4 initiatives in extended pilot phase at 30 months; 0 in production

Phase 1: The Setup (Months 0–6)

The organization launched four AI initiatives simultaneously: a generative AI knowledge base for agents, a predictive fraud detection model, a claims automation agent, and a customer-facing multimodal triage tool. Executive enthusiasm was high. The POC results for all four were impressive. The board approved an accelerated 18-month production roadmap.

EARS diagnostic at Month 6: Layer 1 (Data) scored 35/100. The organization had 14 separate claims data systems, three of which could not be accessed via API. Fraud data was siloed in a regulatory reporting system with no real-time feed. Layer 5 (Governance) had not been started.

Phase 2: The Stall (Months 7–18)

Each initiative hit the same wall: data. The generative AI knowledge base required structured indexing of 40,000 untagged policy documents — a 6-month data remediation project no one had budgeted for. The fraud detection model showed 87% accuracy in testing but fell to 61% in production within 8 weeks due to data distribution shift between the training set (2019–2021) and live claims (2024). The claims automation agent achieved 78% task completion in demos but 34% in the live environment, where edge cases accounted for 41% of actual claims volume.

Failure archetypes at work: POC Trap (all four initiatives), Data Swamp (knowledge base, fraud), Silent Degradation (fraud model), Governance Vacuum (no TRiSM framework, regulatory review triggered in Month 16).

Phase 3: The Remediation (Months 19–30)

A new CTO joined at Month 19 and applied the EARS diagnostic. The decision: stop all four initiatives. Redirect 100% of the remaining budget to Layer 1 and Layer 5 for 9 months before any production deployment.

- Layer 1 (Data): Unified claims data into a single lakehouse with real-time API access. Tagged and indexed 40,000 policy documents. Built a model monitoring pipeline with automatic drift detection.

- Layer 5 (Governance): Stood up the TRiSM framework. Completed regulatory pre-approval for fraud AI. Established HITL review thresholds for claims automation.

By Month 30, the fraud detection model was re-deployed into production with a data governance layer — accuracy held at 91% for the first 6 months. The generative AI knowledge base was launched to 1,200 agents. The claims automation and multimodal tools were deferred to Year 3 (aligned with the EARS strategic build / defer classification in this report).

Lessons Mapped to EARS

Layer 1 (Data): The €12M investment was not wasted on models. It was wasted on deploying models without data readiness.

Layer 3 (Systems): 14 siloed data systems were not a technology problem. It was a data architecture debt that pre-dated AI by 10 years.

Layer 4 (Operations): The fraud model's 26-point accuracy drop in production was not a model failure. It was an operations failure — no drift monitoring, no retraining schedule.

Layer 5 (Governance): The regulatory trigger at Month 16 added 6 months and €800K in unplanned remediation. TRiSM investment at Month 1 would have cost €120K.

The equation: Data Quality was 35/100. No model selection decision could have compensated for that numerator.

The CFO View: How to Evaluate AI Investments

Most AI investment frameworks are written for CTOs. This section is written for CFOs. The financial evaluation of AI differs from traditional IT capital allocation in three critical ways: uncertainty is structural (not just estimation error), returns are nonlinear (front-loaded cost, back-loaded value), and risk compounds across a portfolio of concurrent initiatives.

CapEx vs OpEx Classification

The AI cost structure is predominantly OpEx, not CapEx — a fact that misleads many CFOs into underestimating the total cost of ownership. The typical breakdown:

- CapEx (typically 15–25% of total): Hardware, initial integration engineering, one-time data architecture remediation.
- OpEx (typically 75–85% of total): Inference compute, model, retraining, monitoring infrastructure, compliance overhead, and — most critically — human operations time. Labour accounts for 40–60% of ongoing AI OpEx in most enterprise deployments.

Implication: AI initiatives that appear affordable at year-one CapEx often become cost-prohibitive at year-two OpEx. The CFO should model a 3-year total cost of ownership before approving any AI deployment, with explicit line items for retraining cadence and governance overhead.

Payback Periods by Capability Tier

Capability Tier	Typical Payback	Year-2 ROI	CFO Action
Deploy Now (Score 80–92)	3–14 months	200–600%	Approve with the standard ROI gate
Strategic Build (Score 65–75)	12–24 months	100–300%	Approve with milestone-based release
Experimental (Score <50)	24–48+ months	Highly variable	R&D budget only; no P&L commitment

Risk-Adjusted ROI

Published AI ROI figures are arithmetic means that systematically overstate likely outcomes. The correct CFO metric is expected value, accounting for failure probability:

$$\text{Risk-Adjusted ROI} = (\text{Stated ROI} \times \text{Success Rate}) - (\text{Failure Cost} \times \text{Failure Rate})$$

Example: GenAI project with \$3.70 stated ROI, 65% success rate, \$500K sunk cost on failure

$$\text{Expected Value} = (\$3.70 \times 0.65) - (\$0.5M \times 0.35) = \$2.41 - \$175K$$

Kill Criteria

Every AI initiative approved by the CFO should have explicit kill criteria defined at inception. Recommended thresholds:

- Time gate: If production deployment is not achieved within 2× the projected timeline, initiate kill-or-recommit review.
- Cost gate: If the total project cost exceeds 1.5× the approved budget without a production deployment, freeze all additional spend pending architectural review.
- ROI gate: If the initiative has been in production for 6 months and the measured ROI is below 50% of the projected ROI, classify it as distressed and apply the remediation protocol.
- Data gate: If data readiness scores are below 50/100 at the 3-month mark, pause deployment and redirect budget to Layer 1 remediation.

The most common CFO mistake in AI is not approving bad projects. It is failing to kill them once they are clearly failing.

1. Generative AI for Knowledge Management & Content Operations

- **Deploy Now**

What It Is

Generative AI — primarily in retrieval-augmented generation (RAG) architectures — applied to internal knowledge retrieval, content synthesis, and document operations. The most widely deployed AI capability in 2024–2025.

Adoption Reality

65% of organizations regularly use generative AI (McKinsey, 2024), nearly doubling within ten months. Primary use cases: marketing, sales, product development. Failure rates escalate sharply when moving beyond chat interfaces toward integrated workflows. 30% of GenAI projects will be abandoned after POC by the end of 2025 (Gartner). Root causes: poor data quality, escalating compute costs, and absent governance.

The biggest GenAI risk isn't hallucination. It's a silent degradation at scale.

Economics (ROI + Time-to-Value)

Median time-to-value: 1–4 months. Top performers: \$10.30 ROI per dollar. Average: \$3.70. Vertical AI integrated with proprietary data delivers up to 4x higher ROI than generic tools. Annual cost range: \$50K–\$2M, depending on scale, integration depth, and inference volume.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	78%	Medium	\$\$	\$50K–\$2M / yr	High
Failure Rate	30–42%	Medium			High
Time-to-Value	1–4 Months	Low			High
Readiness Score	85 / 100	Low			High

Executive Decision Block

EXECUTIVE DECISION BLOCK

IF:

- Data readiness score ≥ 60%
- Existing document/knowledge corpus available
- Risk tolerance: Medium

→ THEN: Deploy with RAG architecture; prioritize proprietary data integration over generic tools

Expected Outcome: \$3.70–\$10.30 ROI per dollar; 1–4 month payback; Medium governance risk

Competitive Impact: Cost leadership through 40–60% reduction in knowledge retrieval overhead; speed advantage in content operations.

Real-World Example

A global professional services firm deployed a RAG-based knowledge assistant across 12,000 employees, reducing research time by 40% and achieving full ROI within 3 months. Key success factor: proprietary case data was indexed and governed before launch.

2. AI-Powered Software Engineering & Code Generation • Deploy Now

What It Is

AI-integrated SDLC tools — including coding assistants like GitHub Copilot and Cursor — automating syntax generation, boilerplate creation, code review support, and documentation.

Adoption Reality

In 2024, estimates suggest 30–50% of code is AI-generated in some environments, with GitHub Copilot writing up to 46% of code in certain workflows. Developers accept ~30% of AI suggestions; 88% of accepted code is retained in the final build. Important caveat: junior developers may experience quality regression through over-reliance. Experienced engineers see clear productivity gains.

Accuracy is not the problem. Reliability of AI-generated code in production is.

Economics (ROI + Time-to-Value)

Ramp-up: 2–4 weeks. Year-one ROI is often offset by onboarding and licensing costs (\$50–\$200 per developer per month). Year-two ROI: up to 600% as fixed costs dissipate. Engineering productivity gains: 20–45%. Successful builds up 84%; PR merge rates up 15%.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	65%	Low	\$	\$50K–\$500K / yr	High
Failure Rate	20%	Low			Medium
Time-to-Value	1–3 Months	Low			High
Readiness Score	92 / 100	Low			High

Executive Decision Block

EXECUTIVE DECISION BLOCK

IF:

- Engineering team size ≥ 10 developers
- CI/CD pipeline with automated testing in place
- IP exposure controls implemented

→ **THEN: Deploy with mandatory HITL code review policy; set 6-month quality baseline before expanding access**

Expected Outcome: 20–45% productivity gain; Year-2 ROI 300–600%; Low failure risk

Competitive Impact: Speed-to-market advantage: faster release cycles compress competitor response windows. Margin expansion through reduced senior developer time on boilerplate work.

3. Predictive Maintenance (PdM) & Industrial AI • Deploy Now

What It Is

Machine learning is applied to sensor data and historical failure patterns to predict equipment failures before they occur. Dominant in manufacturing, automotive, and energy sectors.

Adoption Reality

39% of manufacturing organizations have adopted AI for production use cases, with PdM as the primary driver. 95% of adopters report positive ROI. Full payback: 6–14 months. Gartner predicts 70% of large organizations will adopt AI-based maintenance forecasting by 2030.

In industrial AI, a single prevented failure can cover the entire annual platform cost.

Economics (ROI + Time-to-Value)

Documented ROI in high-impact industrial scenarios: 10x–30x within 12–18 months. Unplanned downtime costs mid-size plants \$180K–\$2M per hour. In automotive, one avoided catastrophic failure (\$2.3M/hr), which covers the full annual platform cost. Long-term: 20–40% extension of asset life.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	39%	Low	\$\$\$	\$200K–\$5M / yr	High
Failure Rate	25%	Low			Medium
Time-to-Value	6–14 Months	Low			High
Readiness Score	90 / 100	Low			High

Executive Decision Block

EXECUTIVE DECISION BLOCK

IF:

- Sensor infrastructure on ≥50% of critical assets
- Historical failure data ≥ 2 years
- Asset downtime cost ≥ \$50K/hr

→ **THEN: Deploy targeting high-value assets first; target breakeven within one fiscal quarter**

Expected Outcome: 10–30x ROI (in high-impact industrial scenarios); 6–14 month payback; 95% of adopters positive

Competitive Impact: Risk reduction moat: competitors without PdM face structurally higher maintenance costs and unplanned downtime losses, creating a durable operational cost advantage.

4. Domain-Specific Language Models (DSLMs) • [Deploy Now](#)

What It Is

Language models fine-tuned on industry-specific datasets (e.g., PubMed for healthcare, SEC filings for finance) to achieve higher accuracy and lower inference costs than general-purpose LLMs on specialized tasks.

Adoption Reality

Currently deployed by ~10% of enterprises, but adoption is accelerating. Gartner predicts 50% of enterprise generative AI models will be industry- or function-specific by 2027 (up from 1% in 2023). DSLMs consistently reach 95%+ accuracy on niche tasks, versus 65–75% for general models. Companies fine-tuning domain models see an average 30% accuracy increase over off-the-shelf implementations.

General-purpose LLMs underperform in high-precision environments and create risk in regulated domains.

Economics (ROI + Time-to-Value)

Inference cost advantage: Smaller 7B parameter DSLMs frequently outperform generic 70B models on specific tasks at a fraction of the compute cost. GPT-4 level inference: \$0.02–\$0.12 per token (rapidly declining as of 2025–2026) — unsustainable at enterprise scale. DSLMs offer up to 50% lower development costs vs. building from scratch.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	10%	Medium	\$\$	\$100K–\$1M / yr	Medium
Failure Rate	15%	Low			Medium
Time-to-Value	3–6 Months	Low			High
Readiness Score	80 / 100	Medium			Medium

Executive Decision Block

EXECUTIVE DECISION BLOCK

IF:

- Domain-specific labelled data corpus ≥ 10K documents
- Accuracy requirement ≥ 90% on specialized tasks
- Inference cost reduction required for scale

→ THEN: Fine-tune open-source base model (7B–13B parameters) on proprietary data; avoid building from scratch

Expected Outcome: 30% accuracy gain; 50% inference cost reduction; 15% failure rate

Competitive Impact: Margin expansion through lower inference costs at scale, combined with accuracy advantages that create switching cost moats in regulated industry segments.

5. AI-Enhanced Fraud Detection & Risk Management • Deploy Now

What It Is

AI systems are applied to real-time transaction monitoring, behavioural anomaly detection, and risk scoring. One of the most mature and proven AI applications in the enterprise.

Adoption Reality

73–78% of banks currently use AI for fraud detection. Advanced models report >90% detection accuracy. AI-generated phishing attacks now achieve click-through rates 4.5x higher than traditional methods. Despite progress, 69% of banking professionals believe criminals are currently better at deploying AI offensively than banks are at using it defensively.

Adversarial AI has turned fraud detection into a perpetual arms race. The arms race is accelerating.

Economics (ROI + Time-to-Value)

Global projected savings: £9.6 billion annually by 2026. Early wins within 8–12 weeks. Large-scale risk management: 12–18 month implementation timelines. Annual cost range: \$500K–\$10M, depending on transaction volume and integration complexity.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	78%	Medium	\$\$\$	\$500K–\$10M / yr	High
Failure Rate	25%	Medium			Medium
Time-to-Value	2–4 Months	Low			High
Readiness Score	88 / 100	Medium			High

Executive Decision Block

EXECUTIVE DECISION BLOCK

IF:

- Real-time transaction data pipeline established
- Legacy system APIs documented and accessible
- Adversarial AI monitoring capability exists or is planned

→ **THEN: Shift from reactive monitoring to proactive trust architecture with AI-led remediation**

Expected Outcome: £9.6B industry savings by 2026; 90%+ detection accuracy; 2–4 month TTV

Competitive Impact: Risk reduction moat: superior fraud detection creates a measurable cost-of-loss advantage over competitors operating legacy rule-based systems.

6. Agentic AI (Autonomous Task Agents) • Plan 6–24M

What It Is

AI systems that autonomously plan, reason, and execute multi-step tasks across digital environments without requiring step-by-step human instruction.

Adoption Reality

Approximately 25–30% of enterprises are testing agentic workflows. Early deployments show 40–70% task completion rates for structured workflows. Reliability decay is the dominant failure mode. The 95% confidence threshold required for autonomous operation is rarely achieved in production within the first 18 months.

Agentic AI succeeds at demos and fails at edge cases. Edge cases are 40% of production reality.

⚡ CONTRARIAN VIEW

Agentic AI is overfunded relative to its maturity. The correct investment now is EARS Layer 3–4 readiness, not agent orchestration.

Economics (ROI + Time-to-Value)

Year-one ROI is typically negative. Year-two potential: 200–400% ROI for high-volume, well-scoped workflows. Annual cost range: \$100K–\$2M for off-the-shelf platforms; \$500K–\$5M for custom builds. Hidden orchestration overhead can add \$200K–\$500K annually in poorly structured deployments.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	25–30%	High	\$\$	\$100K–\$5M / yr	Medium
Failure Rate	60–70%	High			Medium
Time-to-Value	6–18 Months	High			Medium
Readiness Score	65 / 100	High			Medium

Executive Decision Block

EXECUTIVE DECISION BLOCK

IF:

- Workflow scope is narrow, well-defined, and bounded
- HITL escalation paths and error recovery exist
- Ops monitoring and audit logging infrastructure in place

→ **THEN: Use off-the-shelf agent platforms for bounded, high-repetition workflows; defer custom orchestration until EARS Layer 4 is mature**

Expected Outcome: 200–400% Year-2 ROI on well-scoped workflows; High initial failure risk on broad scope

Competitive Impact: Speed advantage on high-volume repetitive workflows; cost compression in back-office operations where agent reliability is acceptable.

7. AI-Driven Supply Chain & Demand Forecasting • Plan 6–24M

What It Is

Machine learning is applied to demand sensing, inventory optimization, logistics routing, and supplier risk management.

Adoption Reality

Early adopters report 15–25% inventory cost reductions and 20–35% forecast accuracy improvements. However, multi-tier visibility across global supplier networks requires data standardization efforts, typically taking 12–18 months before AI models can be effectively trained.

Economics (ROI + Time-to-Value)

ROI range: 150–250% in manufacturing deployments. Time-to-value: 12–18 months. Organizations adding external signal data (weather, geopolitical risk, commodity pricing) see 2x higher forecast accuracy. Annual cost range: \$200K–\$5M, depending on tier-1/2/3 supplier integration scope.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	30–35%	Medium	\$\$\$	\$200K–\$5M / yr	Medium
Failure Rate	35%	Medium			Medium
Time-to-Value	12–18 Months	Medium			Medium
Readiness Score	70 / 100	Medium			Medium

Executive Decision Block

EXECUTIVE DECISION BLOCK	
IF:	<ul style="list-style-type: none"> Internal supply chain data standardized across ≥80% of SKUs External signal data integration budget available 18-month timeline acceptable
→ THEN: Build a comprehensive data strategy, including external signals, before model training	
Expected Outcome:	150–250% ROI; 12–18 month payback; 35% failure rate if data is fragmented
Competitive Impact:	Cost leadership through inventory optimization and reduced buffer stock; resilience advantage through predictive disruption detection before competitors react.

8. AI Trust, Risk & Security Management (AI TRiSM) • Plan 6–24M

What It Is

A framework of four technical capability layers — governance, trustworthiness, fairness, and security — providing the operational infrastructure for responsible AI deployment at enterprise scale.

Adoption Reality

44% of organizations have implemented runtime prompt security; 41% have implemented AI red-teaming. Despite this, only 23% of IT leaders are confident in their governance capability. A significant policy-to-practice gap exists: governance on paper without real-time enforcement controls.

You cannot govern what you cannot observe. Most organizations cannot observe their own AI systems.

⚡ CONTRARIAN VIEW

TRiSM will slow down AI adoption before it accelerates it. Near-term (2025–2027), governance operationalization acts as a controlled deceleration. Organisations who delay it face an uncontrolled one.

Economics (ROI + Time-to-Value)

TRiSM operationalization projects to deliver a 50% improvement in AI adoption and user acceptance by 2026. TTV: 4–12 months. Annual cost range: \$100K–\$2M. Risk-adjusted ROI is high given that Gartner predicts a 30% increase in AI legal disputes by 2028.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	44%	Medium	\$\$	\$100K–\$2M / yr	Medium
Failure Rate	40%	High			Medium
Time-to-Value	4–12 Months	Medium			Medium
Readiness Score	75 / 100	Medium			Medium

Executive Decision Block

EXECUTIVE DECISION BLOCK

IF:

- AI model inventory exists or is achievable within 90 days
- Regulatory exposure: GDPR / EU AI Act / SOC 2 in scope
- Existing AI deployments in production

→ **THEN: Inventory all AI assets — models, APIs, agents — and establish baseline risk scores before any new deployment**

Expected Outcome: 50% adoption improvement; 30% legal risk reduction; 4–12 month implementation

Competitive Impact: Regulatory compliance as a competitive moat: organizations with mature TRiSM frameworks will access enterprise customers and regulated markets that competitors without governance cannot serve.

9. Multimodal AI (Vision / Audio / Text Fusion) • Defer 24M+

What It Is

Foundation models trained simultaneously on text, images, audio, and video to enable unified understanding across data modalities. Currently at the Peak of Inflated Expectations.

Adoption Reality

Adoption: ~15%. IDC predicts 80% of production foundation models will be multimodal by 2028 — a 3-year horizon. Current production reliability for complex data fusion remains below enterprise thresholds. Benchmarks show rapid progress (18.8 percentage point improvement on MMMU in a single year), but progress on benchmarks is not the same as production readiness.

Economics (ROI + Time-to-Value)

Most organizations achieve satisfactory returns only after 2–4 years. Hardware and inference costs are prohibitive at current price points. Lack of standardized safety benchmarks creates liability risk in regulated industries. Annual cost range: \$500K–\$10M+ for meaningful production deployments.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	15%	High	\$\$\$	\$500K–\$10M+ / yr	Low
Failure Rate	50%	High			Medium
Time-to-Value	12–24 Months	High			Medium
Readiness Score	45 / 100	High			Low

Executive Decision Block

EXECUTIVE DECISION BLOCK

IF:

- Text-based AI deployments performing well in production
- Hardware and inference budget is constrained
- Hallucination liability exposure in regulated domain

→ THEN: Monitor open-weight multimodal models (Llama 4, Gemini Ultra trajectory); defer production deployment until 2027–2028

Expected Outcome: 2–4 year TTV; 50% failure rate; watch for inference cost curve inflection in 2026–2027

Competitive Impact: Potential future differentiation in complex document processing and visual quality inspection; no near-term competitive necessity for most enterprises.

10. Multiagent Systems (MAS) • Defer 24M+

What It Is

Collections of independent AI agents coordinating to accomplish distributed, branching workflows across different environments.

Adoption Reality

Gartner identifies MAS as a top strategic trend for 2026, but production reality is characterized by early-stage deployment failure rates often cited in the 60–90% range. Most systems are offline within 4–6 months of launch due to orchestration failure, token cost explosion, and feedback loop instability. Current deployment: <5% of enterprises in production.

Multi-agent systems don't fail at launch. They fail at month four, when orchestration debt compounds.

⚡ CONTRARIAN VIEW

Agentic AI is overfunded relative to its maturity. The correct investment is the EARS Foundation, not agent orchestration platforms.

Economics (ROI + Time-to-Value)

Hidden orchestration overhead: \$500K+ annually in poorly structured deployments. Token cost amplification (1.6x–6.2x vs. single-agent) causes finance teams to terminate projects before reaching scale. Time-to-value: 18–36 months in optimistic scenarios. Annual cost range: \$500K–\$5M+, including hidden overhead.

Cost Model

Metric	Value	Risk	Investment	Annual Cost Range	Confidence
Deployment Rate	<5%	High	\$\$\$	\$500K–\$5M+ / yr	Medium
Failure Rate	80–95%	High			High
Time-to-Value	18–36 Months	High			Low
Readiness Score	30 / 100	High			Medium

Executive Decision Block

EXECUTIVE DECISION BLOCK

IF:

- Tolerance for 60–90% pilot failure rate confirmed
- A2A protocol standards not yet established internally
- Agentic AI (single-agent) alternative evaluated and found insufficient

→ **THEN: Defer production investment; allocate small R&D budget to A2A standards; revisit when single-agent deployments are stable in production**

Expected Outcome: 30/100 readiness; 60–90% failure (early-stage deployments); focus on orchestration standards before platform investment

Competitive Impact: Long-term potential for first-mover advantage in complex process automation; no near-term competitive necessity justifies the current failure rate.

Strategic Synthesis & Readiness Matrix

Ranked Readiness Matrix (2024–2025)

Rank	Capability	Score	Verdict	Key Blocker	Risk
1	Software Engineering AI	92	Deploy Now	Junior Developer Quality Gap	Low
2	Predictive Maintenance	90	Deploy Now	Legacy Integration	Low
3	Fraud Detection / Risk	88	Deploy Now	Adversarial AI Evolution	Medium
4	Generative AI (RAG)	85	Deploy Now	Data Quality / Hallucination	Medium
5	DSLMS (Domain-Specific)	80	Deploy Now	Specialized Talent Scarcity	Medium
6	AI TRISM	75	Plan 6–24M	Context Layer Fragmentation	Medium
7	Supply Chain AI	70	Plan 6–24M	Multi-tier Visibility	Medium
8	Agentic AI	65	Plan 6–24M	Reliability Decay	High
9	Multimodal AI	45	Defer 24M+	Prohibitive Hardware Cost	High
10	Multiagent Systems	30	Defer 24M+	Orchestration Overhead	High

The GenAI Divide: Who Wins and Why

The widening gap between the ~5–10% of AI Pacesetters (varies by study) capturing significant EBIT impact and the remaining majority is structural, not technological:

- Winners have proprietary data advantages: They integrate AI with domain-specific datasets that competitors cannot easily replicate.
- Winners invest in operations: 70% of AI resources allocated to people and processes vs. the industry average of 30%.
- Winners govern before they scale: TRISM frameworks are in place before mission-critical workloads are added.
- Losers are trapped in the POC Trap: Impressive demos consuming budget without a production pathway.

The GenAI Divide will widen. The window for catching up is 18–24 months.

AI as a Capital Allocation Problem

AI success is not primarily a technology problem. It is a capital allocation problem under uncertainty. The organizations achieving the highest ROI are not those with the best models — they are those with the most disciplined investment frameworks.

The correct question for CXOs is not: "Which AI capabilities should we deploy?"

The correct question is: "At which EARS layer are we most constrained, and where will capital investment there generate the highest marginal return across our entire AI portfolio?"

The Operating Model Shift

AI is not a tool shift. It is an operating model transformation. The required shifts:

- Organizational structure: Move from functional silos to cross-functional AI product teams with shared data ownership.
- Talent model: Shift from AI specialists to AI-augmented generalists across every function. The scarcest resource is not AI engineers — it is domain experts who can direct AI systems effectively.
- Decision velocity: AI-native organizations make decisions 3–5x faster by delegating pattern recognition to models and reserving human judgment for ambiguous, high-stakes choices.

5 CEO Mandates for 2025

- Mandate 1: Declare a data readiness standard. No AI initiative proceeds without a data quality assessment scoring $\geq 60\%$.
- Mandate 2: Establish TRiSM governance before scaling. Inventory all models, APIs, and agents within 90 days.
- Mandate 3: Reallocate AI budget: 70% to people/process, 30% to technology.
- Mandate 4: Kill all POCs over 6 months old without a production deployment date. Redirect capital to proven capabilities.
- Mandate 5: Build your AI capital allocation framework using the EARS stack as the evaluation lens for every new initiative.

What to Do Monday Morning

- Identify which of the 5 Failure Archetypes is most relevant to your current AI portfolio.
- Map your top 3 active AI projects to the EARS framework — identify the weakest layer.
- Classify each capability investment into the three buckets: Deploy Now, Strategic Build, or Experimental.
- Schedule a TRiSM audit of all existing AI systems within 30 days.
- Set a 6-month ROI checkpoint for every POC currently in flight — kill or commit.

3-Year Outlook (2025–2028)

By 2028, the current top-5 capabilities will be table stakes, not differentiators. The next wave of competitive differentiation will come from organizations that successfully industrialize Agentic AI and Multiagent Systems. The window to build the EARS foundation is now. Organizations that delay this investment will find themselves unable to participate in the next capability wave when it arrives.

Where This Report Could Be Wrong

Overestimation of GenAI ROI: Published figures skew toward early adopters and may not reflect average enterprise outcomes at scale.

Underestimation of regulatory impact: EU AI Act timelines could compress deployment windows faster than current projections.

Cost curve uncertainty: Inference costs are declining ~50–80% per year. Cost barriers to multimodal and MAS may be overstated for 2027+ deployments.

Organizational capability assumption: For enterprises in early digital transformation stages, all timelines should be extended by 6–12 months.

If You Are a CXO Reading This...

This report is designed to be useful in a boardroom. The following section converts the research into a self-diagnostic. Use it before your next AI investment decision, your next board update, or your next conversation with a technology vendor.

3 Questions to Ask Yourself Right Now

Question 1: What is our EARS Layer 1 score?

If you cannot answer this within 60 seconds — with a number, not a narrative — your organization has a data readiness problem.

The question is not “Do we have data?” Every organization has data.

The question is: “Is our data accessible, labelled, governed, and pipeline-ready for the specific AI capabilities we are trying to deploy?”

If the honest answer is no, no model selection decision will compensate. Stop and fix Layer 1 first.

Question 2: Which of our active AI initiatives is in the POC Trap right now?

Every organization with more than 2 active AI initiatives has at least one in the POC Trap.

The diagnostic: identify every AI initiative that has been in ‘pilot’ or ‘evaluation’ status for more than 6 months.

For each one, ask: what specific, dated production deployment milestone exists?

If the answer is “we’re still evaluating,” that initiative is consuming budget with no ROI pathway. Kill it or set a 90-day commit-or-close deadline.

Question 3: If AI succeeds at the scale we are claiming, are we operationally ready?

Most AI business cases are written for launch day. Almost no model in Month 18.

The question is: if this capability achieves the projected adoption and usage volume, do we have the monitoring infrastructure, retraining budget, governance controls, and talent to sustain it?

If the answer is ‘we’ll deal with that when we get there,’ you are pre-building a Silent Degradation or Governance Vacuum failure.

Model Month 18 explicitly before approving any initiative.

3 Signals You Are Failing at AI — Right Now

- Signal 1: Your AI team cannot tell you the current production accuracy of your deployed models within 24 hours. This is not a data problem. It is an operations maturity problem (EARS Layer 4 score: <40). If you cannot measure it, you cannot manage it — and you will not know when it fails.
- Signal 2: Your AI roadmap contains more than 3 concurrent initiatives with the same resource pool. Parallelism in AI is a trap. Each additional concurrent initiative reduces the depth of data readiness, governance, and operational infrastructure available per initiative. The organizations achieving the highest ROI run 1–2 deep initiatives, not 6–8 shallow ones.
- Signal 3: Your AI investment thesis is technology-led, not outcome-led. If the primary justification for an AI initiative is “we need to stay current with GenAI” or ‘competitors are deploying Agentic AI,’ you are making a fear-of-missing-out capital allocation decision. The AI Execution Equation predicts the outcome: low numerator variables (data quality, operational maturity, adoption depth) produce negative ROI regardless of model quality.

The gap between AI leaders and laggards is not technical. It is diagnostic. Leaders know exactly where they are failing before they spend another dollar.

Works Cited

1. The State of AI in Early 2024 — McKinsey: [mckinsey.com/capabilities/quantumblack](https://www.mckinsey.com/capabilities/quantumblack)
2. Gartner: 30% of GenAI Projects Will Be Abandoned After POC by the End of 2025
3. Why 95% of Corporate AI Projects Fail — ComplexDiscovery / MIT 2025 Study
4. Stanford HAI: 2025 AI Index Report — hai.stanford.edu/ai-index/2025-ai-index-report
5. Gartner: Survey Finds GenAI Most Frequently Deployed AI Solution (May 2024)
6. Gartner: 70% of Large Organizations to Adopt AI-Based Supply Chain Forecasting by 2030
7. Gartner: AI Regulatory Violations Will Increase Legal Disputes 30% by 2028
8. IDC: 80% of AI Foundation Models Will Be Multimodal by 2028
9. Gartner: Top Strategic Technology Trends for 2026 (October 2025)
10. Business Wire: 75% of Enterprises Report Double-Digit AI Failure Rates (March 2026)
11. Gartner: Lack of AI-Ready Data Puts AI Projects at Risk (February 2025)
12. BioCatch: 2024 AI Fraud Financial Crime Survey
13. AI Supply Chain Playbook: 150–250% ROI with AI Agents in 2026 — aimagicx.com
14. Gartner: 60% of Supply Chain Disruptions Resolved Without Human Intervention by 2031
15. Fast Company: AI ROI Is on the Rise — [fastcompany.com/91479632](https://www.fastcompany.com/91479632)
16. Silent Failure at Scale: The AI Risk That Compounds Before Anyone Notices — beam.ai
17. Gartner Hype Cycle: Top AI Innovations in 2025 (August 2025)
18. WalkMe: State of Enterprise AI Adoption 2025
19. Fullview AI: 200+ AI Statistics & Trends for 2025
20. ResearchGate: Why ~95% of Enterprise AI Initiatives Fail to Deliver Financial ROI