Moving from investment to impact with AI projects

Why operational discipline, governance, and workflow redesign determine whether AI investments create measurable enterprise value.

By
Martin Druda
Cloud Strategy and AI Solutions,
e-Core North America

OVERVIEW

Executive Summary

AI investment continues to accelerate across industries, yet many organizations still struggle to demonstrate measurable business impact from their AI initiatives. The gap between experimentation and sustainable ROI remains significant.

Some organizations are reporting measurable gains in efficiency, productivity, and operational performance. Others remain stuck in disconnected pilots, unclear business cases, and expensive experimentation with little path to scale.

Access to AI capabilities is no longer the bottleneck. Most companies have full access to AI technologies, to advanced models, automation platforms, and rapidly evolving agentic capabilities. Now, the challenge is translating those into measurable outcomes that improve revenue, margins, efficiency, or customer experience.

The organizations who excel on that are seeing strongest returns. They share several characteristics, such as beginning with clearly defined business challenges rather than technology-first initiatives. They scope projects narrowly enough to prove value quickly. They establish measurable baselines before deployment. They redesign workflows instead of layering AI onto broken processes. And they treat governance, security, and human oversight as foundational components of implementation rather than secondary concerns.

At the same time, agentic AI is accelerating the urgency of these decisions. Systems capable of reasoning, planning, using tools, and executing multi-steps tasks are beginning to move from experimentation into production environments. That shift introduces new opportunities for operational scale, but also new governance, security, and accountability challenges.

The organizations achieving measurable ROI from AI are not necessarily investing the most aggressively.

They are investing with the clearest alignment between business objectives, implementation strategy, governance, and operational execution.

74% of executives report achieving measurable ROI from AI initiatives within the first year of deployment.[1]

TAKEAWAYS

What you'll get from this paper

This paper is designed for enterprise leaders navigating the growing pressure to operationalize AI investments responsibly, efficiently, and at scale.

How to evaluate AI opportunities before investment

Learn how to distinguish between AI opportunities, traditional automation needs, and workflow redesign challenges before committing resources.

Why most AI pilots fail to reach production

Understand the operational, governance, and organizational gaps preventing measurable ROI.

Where agentic AI creates measurable business value

Explore the characteristics of high-value agentic AI use cases and the operational environments where they succeed.

How governance becomes a competitive advantage

Understand how trust architecture, human oversight, auditability, and access governance shape long-term adoption.

How leading organizations move from experimentation to operational impact

See how disciplined scoping, workflow redesign, and operational alignment change outcomes.

What executives should rethink immediately

Prepare for the next phase of enterprise AI adoption with clearer operational priorities and governance strategies.

OUTCOMES

The investment reality: Your mileage may vary

Generative AI investment is accelerating rapidly across industries, creating both opportunity and pressure for enterprise leadership teams. Yet while some organizations are beginning to report measurable operational gains, many others remain trapped in disconnected pilots, unclear ROI, and fragmented strategies.

The data reveals a gap between companies capturing real value and those accumulating expensive lessons.

KEY STAT:

74% of executives report achieving measurable ROI from AI initiatives within the first year of deployment.[1]

KEY STAT:

95% of AI pilot programs still fail to demonstrate meaningful P&L impact at production scale.[2]

KEY STAT:

31.5% year-over-year increase in agentic AI becoming a top enterprise IT investment priority.[3]

The pre-investment decision framework

Before organizations commit to any AI investment, the most important decisions are often operational rather than technical. The highest-performing initiatives begin with realistic assessments of organizational readiness.

Executives and their teams should work through an organized set of questions as a guide that filters high-ROI initiatives from expensive experiments.

1.

Start with the business problem, not the technology

While it sounds foundational, it doesn't always happen. There is pressure to deploy AI from multiple fronts: boards, competitors, vendor briefings. The pressure creates a sense of urgency toward technology-first decisions.

The result is what consultants now call "solution shopping": acquiring capabilities before defining the problem they are meant to solve. It is essential to understand three critical components of the decision framework:

- State the business problem in business language, not technical language.
- Define what a solved problem looks like before selecting a solution.
- Identify a named business owner accountable for outcomes.

2.
Is AI actually the right tool?

One of the most valuable questions an organization can ask is whether a selected use case requires AI at all. There are a multitude of automation solutions, and AI (particularly generative and agentic AI) is not always the appropriate solution.

- Is this use case deterministic? If so, traditional automation is likely a better fit
- What happens when the system is wrong and how recoverable is the error?
- Does this process require reasoning or just retrieval?

What does "Deterministic" mean?

Deterministic processes follow fixed rules to produce predictable outcomes from identical inputs. These workflows are generally better served by traditional automation than AI.

Examples of deterministic tasks:

- routing invoices based on fixed criteria
- sending approval notifications
- validating structured form fields
- applying predefined business rules

AI becomes more valuable when the workflow involves:

- reasoning
- ambiguity
- unstructured data
- contextual decision-making
- natural language interaction

Practical question for leaders:

"If the process already follows fixed rules with predictable outcomes, do we need AI at all?"

3.

Iterate: start small, prove value, scale

The companies achieving the highest returns from AI have something in common: they select a limited, well-defined use case, instrument it for measurement, prove ROI, and then use that proof to build internal capability and organizational confidence before iterating.

51% of organizations now achieve a production-ready product in three to six months.1

- Scope the use case narrowly enough to prove value within three to six months
- Define success criteria in advance: what would cause us to scale? What would cause us to stop?
- A pilot without exit criteria is an experiment with no end date and no budget accountability

What are exit criteria?

Exit criteria define the measurable conditions that determine whether an AI pilot should scale, pause, or stop. Without them, pilots often become open-ended experiments with unclear ROI and no accountability.

Examples include:

- cost reduction targets
- productivity gains
- response time improvements
- accuracy thresholds
- adoption goals

Practical question:

"What measurable outcome would justify scaling this initiative and what result would tell us to stop?"

AI Decision Framework

Flow:

Business Problem → Does this require AI? → Deterministic or reasoning-based? → Data readiness assessment → Workflow impact analysis → Pilot scope definition → Success criteria established → Production readiness

DEPLOYMENT

Agentic AI strategy: From pilot to production

Agentic AI introduces a new operational model for enterprise technology. Unlike traditional automation systems, agentic architectures can reason across workflows, use tools dynamically, execute multi-step tasks, and operate with limited human intervention. That shift changes not only implementation strategy, but governance, security, and operational accountability itself.

Thinking about Agentic, there is the capability to autonomously plan, reason, use tools, and execute multi-step tasks with minimal human oversight. This represents a significant shift in corporate technology. The adoption data shows this: 52% of leaders report their organizations are using AI agents in production, and 39% have deployed more than 10 agents.

Among companies that have committed at least half their AI budget to agentic systems, 88% report positive ROI.1

Agentic AI Market Projection

$5.25B → $199B

projected ten-year market growth by 2034

43.84% CAGR

Source: Grand View Research/GlobeNewswire, 2025.

The right use cases for Agentic AI

High ROI agentic implementations have common characteristics:

They involve multi-step workflows that previously required human coordination

They operate on unstructured or variable inputs

They benefit from the ability to use multiple tools or data sources in sequence and the cost of iteration (getting an answer slightly wrong and correcting it) is low relative to the cost of the manual process it replaced.

Google Cloud's 2025 research identifies leading use cases: customer service (49%), marketing (46%), security operations (46%), tech support (45%), software development

WSJ highlights the results of early adopters:

- BNY deployed ~100 AI agents as "digital employees" with distinct logins, email access, and human managers; one agent scans and patches code vulnerabilities
- Walmart's Trend-to-Product agent shortened the fashion production timeline by up to 18 weeks

Both examples share limited scope, defined authority, measurable outcomes, and human oversight built in from the start.

CASE EXPERIENCE

I've seen this pattern across industries and at different scales. We worked with a real estate data company collecting over 20 million property listings daily from more than 11,000 websites. They faced a different version of the same problem: 100 scraping scripts failed every day, and each manual repair took approximately eight hours to complete.

We delivered an AI-powered system built on Amazon Bedrock that automated both script generation and repair, cutting that cycle from eight hours to 30 minutes and reducing the operational cost of building 200 new scripts by 78%. The ROI came from a well-scoped problem, a measurable baseline, and a clear definition of success established before any build began.

GOVERNANCE

Guardrails, hallucinations, and the trust architecture

The largest barriers to enterprise AI adoption are increasingly related to trust, accuracy, governance, and operational accountability rather than model capability itself.

In the MIT NANDA study, ~33% of the professionals cited hallucinations and data reliability as their primary concern, outweighing cost or unclear ROI.[2] To build trust into the deployment architecture from the start, organizations must:

- Define when and how model outputs require human validation before action is taken
- Address key design questions:
- Where does a hallucination in this workflow create irreversible harm?
- How are we grounding the model in authoritative, current data?
- How do we audit what the agent did and why?
- What happens when the agent encounters an out-of-scope request?

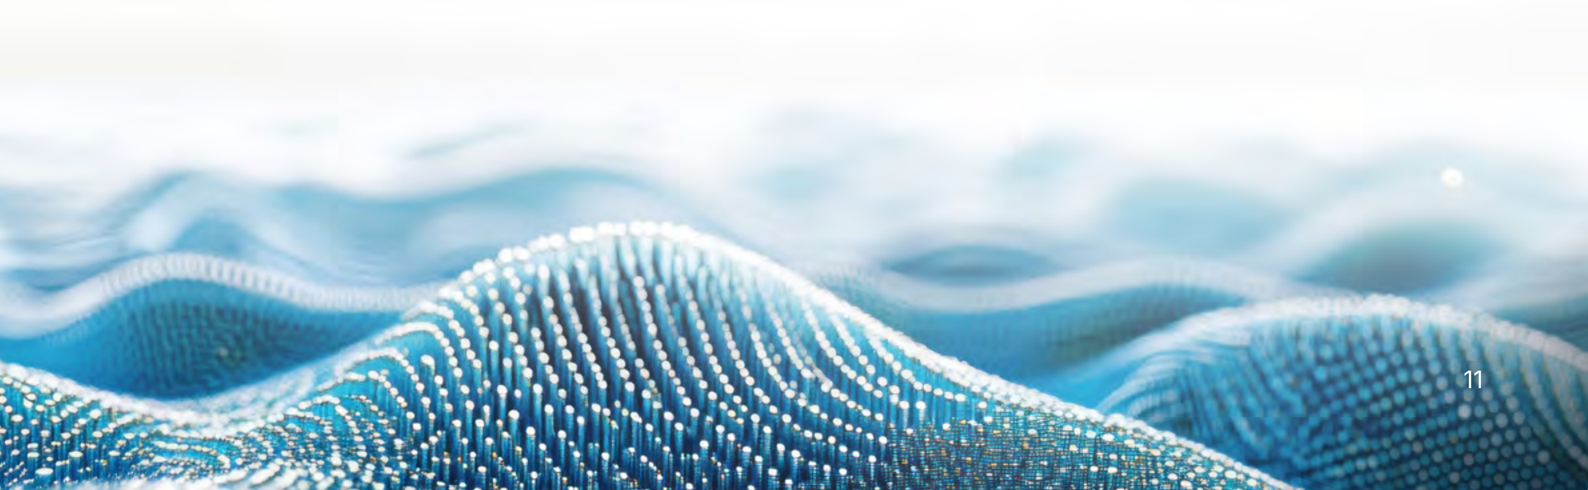

COMPLIANCE

Security and access governance

As AI agents begin to take actions such as: creating records, sending communications, accessing systems, and executing transactions, the security and access model governing the agents becomes a risk concern. The traditional identity and access management (IAM) framework was designed for human users and service accounts. Agentic AI introduces a new class of identity that requires a new governance approach.

OpenClaw case study

OpenClaw, an open-source platform that allows organizations to deploy autonomous AI agents capable of interacting with systems, tools, and enterprise workflows, gained massive adoption shortly after launch because of how quickly teams could operationalize agentic AI capabilities.

The platform reached 180,000 GitHub stars in weeks, but security researchers also identified tens of thousands of exposed instances leaking credentials and sensitive access data.[7]

The issue was not OpenClaw itself, but the speed at which organizations deployed powerful autonomous agents without governance, visibility, or proper security controls in place.

- Employees deployed a powerful agentic tool to corporate devices without IT visibility; traditional endpoint, network, and identity tools could not see it [7]
- Cisco called it a groundbreaking capability, but also a significant risk from a security perspective [7]
- At GTC 2026, Nvidia's Jensen Huang showcased NemoClaw: "every company needs an OpenClaw strategy; this is the new computer" [5]

Key Governance Requirements

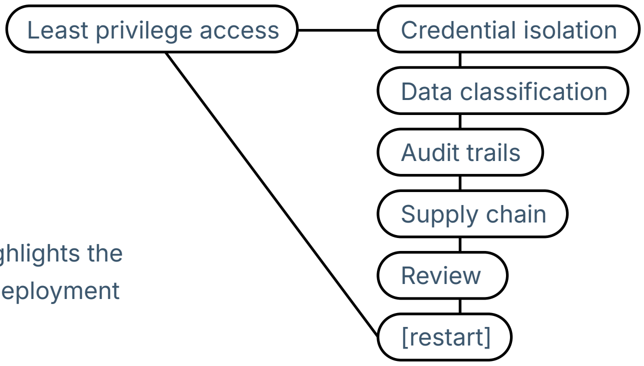

OpenClaw is not the threat; it highlights the governance gap every agentic deployment will face.

Security design requirements for agentic AI:

- Principle of least privilege: agents get only the permissions their task requires
- Credential isolation: each agent identity is individually auditable; no shared credentials across workflows
- Data residency and classification: explicit policies on what agents can access, transmit, and store
- Audit trails: log every agent action for compliance and incident review
- Supply chain risk: third-party AI components and integrations require the same scrutiny as any external software dependency

37% of executives now cite data privacy and security as their top consideration when selecting AI providers[1]

STRATEGY

Build, Buy, or Partner?

The question is no longer whether organizations should adopt AI capabilities. The question is how they should operationalize them responsibly and efficiently.

The MIT NANDA report analyzes the strategic benefits and potential pitfalls of the three approaches: AI initiatives that involved purchasing solutions from specialized vendors and building partnerships succeeded approximately 67% of the time. Internal builds succeeded only 33% as often. Yet many companies were attempting to build proprietary AI systems.

This is not an argument against building; it is an argument for honesty about organizational capability

Decision factors: differentiation value, AI engineering maturity, time-to-value, total cost of build vs. buy, long-term maintenance

For most organizations, partner-led implementation with vendor-provided infrastructure and internally-owned customization is the right balance

The 67% partnership success rate versus one-third for internal builds highlights a consistent gap: organizations often underestimate the complexity of AI engineering and the operational work required after deployment, including monitoring, guardrails, retraining, and continuous improvement. Partner-led implementations combine engineering depth with operational experience, which helps explain their higher success rates.

FACTOR	INTERNAL BUILDING	VENDOR SOLUTION	STRATEGIC PARTNERSHIP
Time-to-value	Slow	Fast	Moderate
Customization	High	Moderate	High
Operational burden	High	Lower	Shared
Governance flexibility	High	Limited	High
AI engineering maturity required	Very high	Low	Moderate
Long-term adaptability	High	Moderate	High

CASE STUDY

What does a productive partnership look like in practice?

The strongest engagements start with a structured discovery process: mapping existing workflows, establishing current-state performance baselines, and identifying where AI creates genuine leverage rather than where it simply appears to create opportunity.

A good example is a fintech accounting company we worked with that wanted to reduce the time accountants spent reviewing dense tax reports. We enabled natural language querying across 15-to-25-page reports, eliminating manual search across portfolios of 100 to 150 companies per professional.

Using AWS Bedrock, structured prompt engineering, and a multi-phase architecture for query generation and summarization, we reduced response times to under 20 seconds.

The client brought domain expertise and proprietary data; we brought the AI engineering depth required to deliver the solution reliably at production scale.

That division of responsibility is what strong partnerships look like: the organization owns business outcomes and strategic direction, while the partner provides the technical execution and operational discipline needed to move from pilot to production.

EVOLUTION

Business-Aligned Modernization

Organizations achieving measurable AI ROI are not treating AI as isolated technology deployments. They are treating it as operational modernization tied directly to business outcomes, workflow redesign, and organizational change.

[box para destaque]
AI initiatives are not just IT projects. They are business transformation programs that happen to have a significant technology component. The companies that achieve ROI treat them accordingly, with executive sponsorship at the business level, cross-functional governance, and outcomes defined in business language at the start.

CASE STUDY

Workflow redesign is critical

The MIT finding that 95% of AI pilots fail to show ROI is partly a technology failure, but it is primarily a workflow design failure. According to the study, companies attempted to bolt AI onto existing processes rather than redesigning those processes around AI capabilities.[2]

- Winning organizations rebuild processes and re-platform content and knowledge; they do not simply add AI to existing steps

- The organizations that did the workflow work are seeing the returns; those that skipped it are running expensive pilots with no clear path to production

The difference between bolting AI onto existing processes and redesigning processes around AI capabilities is visible in outcomes. When e-Core partnered with a large on-demand delivery platform processing over **120 million orders monthly**, the engagement didn't begin with model selection.

We started with a complete mapping of the existing operation: escalation flows, alert configurations, response bottlenecks, and SLAs by area. That groundwork revealed that the fundamental problem wasn't response speed in isolation; it was the structure of the support model itself.

The team designed an entirely new operating model around that insight, a proactive Command Center built on standardized runbooks, AI-driven alert triage, and correlation with complete end-to-end follow-through on every incident. **Mean time to respond dropped from 8 minutes to 47 seconds.**

Mean time to resolve dropped from **180 minutes to 55 minutes.** False-positive alerts were **reduced by 88%**.

The results didn't come from the AI alone. They came from redesigning the workflow around it.

FOUNDATION

Data Readiness is a Prerequisite

PwC research indicates that 61% of organizations report their data sources are not ready for AI implementation.8 This is the primary reason why technically capable AI systems fail in enterprise environments. **An AI agent is only as effective as the data it operates on.**

Domains to assess before deployment: completeness and accuracy in the target domain; governance and lineage documentation; accessibility for AI consumption; sensitivity classification; quality assurance processes

Isolated systems create a real impact on every AI use case, especially in work integration, data reconciliation, latency, and governance complexity

Address data and integration readiness before deployment, not after

Data readiness challenges often surface before deployment when organizations examine their data carefully.

A logistics risk management firm tracking millions of trips discovered that claims represented fewer than 2% of events, creating a heavily imbalanced dataset. Without correction, the model would have learned to predict the majority class while missing the high-risk cases that mattered most.

By addressing the imbalance before algorithm selection, the resulting model achieved a **94% AUC and correctly identified 88.5% of at-risk trips**. Organizations that skip this assessment step often discover these issues only after deployment, when models that performed well in testing begin producing unreliable real-world results.

To put it another way: "**About 9 out of 10 high-risk trips were correctly identified**". Organizations that skip the data assessment step tend to discover this problem after deployment. That's when a model that performed well in testing begins producing unreliable results against real-world inputs where the data distribution differs from what the training set assumed.

61% of organizations report that data sources are not ready for AI implementation.

CULTURE

Talent, change management, and the human side

Operational transformation does not happen through technology adoption alone. AI initiatives succeed when organizations redesign workflows, retrain teams, and align operational culture around new ways of working.

Inadequate infrastructure is a leading cause of AI project failure, but human factors account for a similar share of failures. Leaders see skills gaps, workforce resistance, and cultural barriers as contributing factors in underperforming AI initiatives.

In Google Cloud's 2025 research, 39% of organizations reporting productivity gains from AI initiatives said productivity at least doubled.[1]

Invest in retraining, redesign roles around augmented capability, and treat change management as a first-order deliverable

The Futurum Group's 2026 research: organizations capturing the greatest returns are winning on people and process investment; not model selection or vendor choice.[3]

Executive Ownership and Governance

One of the significant findings in the McKinsey State of AI research is the role of executive ownership. 62% of larger organizations have established a dedicated team to drive AI adoption, but only 28% report their CEO as responsible for AI governance, and only 17% demonstrate board-level oversight.

AI initiatives without visible senior leadership ownership tend to stall at the pilot stage because they lack the authority to bridge functional boundaries, secure budget, and drive process change.[9]

Enterprise AI Governance Structure

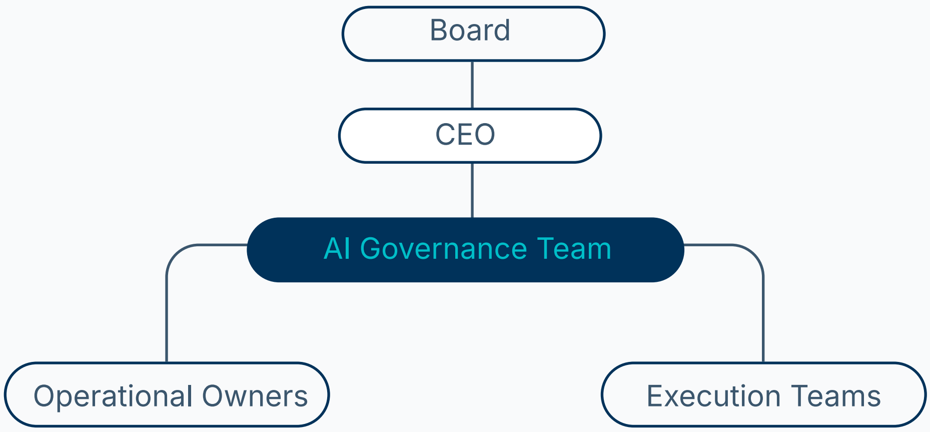

CONTEXT

Regulatory landscape and competitive urgency

Two external factors influence the strategic timeline for AI investment: an evolving regulatory environment and a shrinking window for competitive differentiation.

The Regulatory Dimension

The EU AI Act is currently the most comprehensive AI regulatory framework in effect. It established risk-tiered obligations for AI systems used in high-risk domains, including: healthcare, financial services, employment, and critical infrastructure. US federal and state-level AI legislation is advancing on multiple fronts, with sector-specific regulators: OCC, CFPB, HHS, and others issuing increasingly detailed rules on AI use in regulated activities.

- The regulatory question is not whether compliance obligations will apply, as they will
- Build governance infrastructure now that accommodates obligations as they mature
- The same architecture that satisfies regulatory requirements also produces trustworthy AI systems; compliance and performance have synergy.

The Competitive Urgency

The advantage of early, structured AI adoption is visible in the results. Companies that have moved from pilot to production with proper infrastructure are seeing productivity gains, revenue growth, and customer experience improvements that their slower-moving competitors are not achieving.

- Agentic AI surged 31.5% as the top technology priority in a single year; the market has moved from exploration to competition[3]
- Organizations establishing agentic capability today are building operational expertise, process knowledge, and integration architecture that will be difficult to replicate quickly
- The decisions made in the next 12 months will create compounding advantages or disadvantages

CHECKLIST

Executive decision checklist

Use this checklist before committing an investment to any AI initiative.

Business Alignment

- [ ] Can we define the business problem without relying on technical language?
- [ ] Is there a named business owner accountable for outcomes, not just IT ownership?
- [ ] Do we have measurable baseline metrics for the current process or workflow?
- [ ] Have we defined what success looks like before implementation begins?
- [ ] Have we established clear criteria for scaling, pausing, or stopping the initiative?
- [ ] Does this initiative connect to measurable business outcomes such as revenue growth, margin improvement, operational efficiency, or customer experience?

Use Case Validation

- [ ] Have we confirmed AI is the right solution for this problem?
- [ ] Would traditional automation or workflow redesign solve the problem more effectively?
- [ ] Does the use case require reasoning, decision-making, or unstructured data handling?
- [ ] Have we identified where the system can fail and what the business impact would be?
- [ ] Is the use case scoped narrowly enough to demonstrate value within three to six months?
- [ ] Have we realistically assessed build vs. buy vs. partner based on our internal capabilities?

Data & Infrastructure Readiness

- [ ] Is the required data accurate, current, complete, and accessible?
- [ ] Do we understand the structure and distribution of the data being used?
- [ ] Have we identified integration dependencies between systems and platforms?
- [ ] Is the data governed, classified, and documented appropriately?
- [ ] Do we have lineage, auditability, and quality assurance processes in place?
- [ ] Are our systems capable of supporting AI workflows without significant reconciliation overhead?

Security & Governance

- [ ] Have we defined least-privilege access policies for AI agents and systems?
- [ ] Are credentials isolated, auditable, and secured appropriately?
- [ ] Do we have logging and traceability for all AI-driven actions?
- [ ] Have we defined where human oversight is required within the workflow?
- [ ] Do we understand the regulatory and compliance implications of this use case?
- [ ] Do we have policies for handling hallucinations, incorrect outputs, and out-of-scope requests?
- [ ] Are third-party AI tools and integrations reviewed with the same scrutiny as other enterprise software dependencies?

Organizational Readiness

- [ ] Is there visible executive sponsorship for the initiative?
- [ ] Have we aligned business, operations, IT, security, and compliance stakeholders?
- [ ] Do we have a workforce transition or retraining plan where necessary?
- [ ] Are teams prepared to redesign workflows rather than layer AI onto existing processes?
- [ ] Is governance designed to accelerate responsible adoption instead of slowing execution?
- [ ] Have we communicated the expected ROI and business rationale clearly to leadership and stakeholders?
- [ ] Do we have operational ownership for maintaining, monitoring, and improving the system after deployment?

EXECUTION

Discipline is the differentiator

The data is clear: companies that realize ROI from AI have discipline.

Discipline to define problems before selecting solutions, to establish baselines before claiming improvement, to govern with purpose rather than caution, and to treat workforce and workflow transformation as primary challenges rather than secondary considerations.

The gap between experimentation and operational transformation is widening. Organizations that establish disciplined adoption models now are building operational knowledge, governance maturity, and workflow infrastructure that will compound over time.

Futurum's 2026 research put it bluntly: AI investments must now connect to revenue growth or margin improvement.3

Productivity arguments alone no longer win approval. The organizations that established those connections through disciplined scoping, data investment, and governed deployment are already leveraging that advantage.

The strategic opportunity remains open, but the next phase of enterprise AI adoption will reward organizations that approach implementation with clarity rather than urgency.

The question is no longer whether AI will reshape enterprise operations. It will. The question is whether organizations will approach that transformation deliberately enough to create a sustainable business impact from it.

BIBLIOGRAPHY

Sources and Research References

Superscript numbers throughout this document correspond to the sources below.

1. **Google Cloud / National Research Group**. ROI of AI 2025. Survey of 3,466 senior leaders across 24 countries. September 2025.
2. **MIT NANDA Initiative / Aditya Challapally.** The GenAI Divide: State of AI in Business 2025. Based on 150 executive interviews, 350 employee surveys, and 300 public AI deployment analyses. 2025.
3. **Futurum Group.** 1H 2026 Enterprise Software Decision Maker Survey. Survey of 830 global IT decision-makers. February 2026.
4. **Globe Newswire / Grand View Research**. Agentic AI Market Sizing and CAGR Projections, 2024-2034. Aggregated statistics. 2025.
5. **Rosenbush, Steven. Wall Street Journal**. Companies Begin to See a Return on AI Agents. November 12, 2025. Subscription required.
6. **Bousquette, Isabelle and Rosenbush, Steven. Wall Street Journal**. Nvidia Software Aims to Bring OpenClaw to the Enterprise. March 16, 2026. Subscription required.
7. **CrowdStrike / Cisco / Security Research Community**. OpenClaw Security Assessments. Multiple sources: CrowdStrike blog Feb. 2026; Cisco blog Jan. 2026; VentureBeat Jan. 2026.
8. **PwC.** Technology Vision 2025. AI organizational readiness and data readiness research.
9. **McKinsey & Company**. The State of AI 2025. Survey of 1,993 participants across 105 nations. November 2025.

e-core

Engineering to scale.

Built to last.