

Enterprise AI Adoption Decisioning

A practical executive guide to value, risk, and the AI ecosystem you must build first

“The bottleneck is never the bot. The AI agents are ready. Most organizations aren’t.”

Board / CEO Briefing

Tunde Kehinde

Managing Partner, 5054 Holdings

Cyril Simone

AI, Machine Learning & Quantum Computing

July 2026 Edition

Updated quarterly. This paper reflects the state of enterprise AI as of Q3 2026.

Contents

1. Executive Summary
2. What Changed Since the February 2026 Edition (Q3 2026 Update)
3. The Executive Decision Framework: AI Adoption + Agent Autonomy
4. The Business Case for “AI in General”
5. Organizational Ownership and Accountability
6. Executive Support: The Air Cover That Makes or Breaks AI Programs
7. Talent and Operating Model
8. When (and Why) to Move from Copilots to Agents
9. The Vendor Decision
10. The Required AI Ecosystem (What You Must Build)
11. Risk, Compliance, and Non-Negotiable Controls
12. Understanding the AI Cost Model
13. Roadmap and Funding Logic (12-18 months)
14. The Transformation Gap: What Enterprises Actually Face
15. Executive Appendix: Toolkit
16. About the Authors
17. Bibliography

Executive Summary

AI adoption is not a binary decision. It is a maturity ladder. The smart move is to incorporate AI in phases, starting with measurable, low-blast-radius use cases, then graduating to agents only after you have built the ecosystem that keeps them safe.

Self-learning and self-modifying agents are the top of the ladder, not the starting point. They can create step-change productivity, but they also create step-change risk. If your organization does not have strong data governance, auditability, tool controls, evaluation discipline, and incident response, agents will eventually do something expensive and embarrassing, at scale.

The most important thing to understand before reading further: AI is not a tooling purchase. It is an operating model change. Buying a copilot license and giving it to your team is not AI adoption. Adoption means redesigning workflows, redefining roles, building new evaluation and monitoring disciplines, and committing to a sustained learning curve. Organizations that treat AI as a procurement decision rather than a transformation program consistently underperform, regardless of how much they spend.

Executives commonly fall into several traps when approaching AI. They assume that deploying a copilot will automatically produce ROI without integrating it into existing workflows and retraining teams to use it effectively. They ignore the need for evaluation and monitoring infrastructure until a high-profile failure forces a reaction. They underestimate the organizational politics: who owns AI, who funds it, who is accountable when it fails, and who has the authority to push adoption through resistant business units. And they skip the maturity ladder entirely, jumping to autonomous agents before the foundational capabilities are in place, which creates risk that is disproportionate to the value delivered.

The first half of 2026 validated this thesis rather than weakening it. Agentic AI now sits at the top of the hype curve: Gartner's first dedicated Hype Cycle for Agentic AI (April 2026) places the category at the Peak of Inflated Expectations, and its 2026 CIO and Technology Executive Survey found that only 17% of organizations have deployed AI agents to date, even though more than 60% expect to within two years. Gartner continues to project that over 40% of agentic AI projects will be canceled by the end of 2027 due to unclear value, rising cost, and weak governance. Meanwhile, the market has coined a term for the gap between marketing and reality: "agent-washing." The bottleneck remains organizational readiness, not model capability.

This paper provides a structured decision framework, a maturity model for graduating through levels of AI autonomy, and the operational ecosystem you must build at each stage. It also addresses the questions that most commonly stall enterprise adoption: who owns AI inside the organization, what support they need from the top, what team is required, how to evaluate vendors without creating lock-in, and how to think about the cost model.

Evidence anchors (selected external findings used to ground this paper):

- Productivity gains are real but uneven: field experiments show ~14% average lift in customer support with larger gains for less-experienced workers (Brynjolfsson et al., NBER 2023).
- Controlled studies show ~40% faster completion on professional writing tasks with generative AI (Noy & Zhang, 2023) and faster completion on software tasks with coding copilots (Peng et al., 2023).

- Adoption is broadening but still shallow: 19.8% of US businesses reported using AI in a business function as of May 2026 (32% employment-weighted), and 57% of adopters use AI in three or fewer business functions (US Census Bureau BTOS, 2026).
- Agent deployment lags agent ambition: only 17% of organizations have deployed AI agents, versus 60%+ that plan to within two years; Gartner projects 40%+ of agentic AI projects will be canceled by end-2027 (Gartner, 2026).
- Agentic systems amplify security risk: OWASP guidance elevates prompt injection and excessive agency/tool abuse; primary research demonstrates indirect prompt injection against real LLM-integrated applications (especially when RAG pulls untrusted content). A 2026 Gartner survey found 74% of IT application leaders view AI agents as a new attack vector, while only 13% strongly agree they have the right governance structures in place.
- Governance expectations are converging even as timelines shift: NIST AI RMF + NIST GenAI Profile, ISO/IEC 42001, and the EU AI Act all emphasize lifecycle risk management, testing, logging, human oversight, and post-deployment monitoring. The EU's Digital Omnibus on AI (agreed May 2026) deferred high-risk system obligations to December 2027 / August 2028 but left the August 2026 transparency obligations in place.

What Changed Since the February 2026 Edition (Q3 2026 Update)

This paper is updated quarterly. Five developments in the first half of 2026 are material to the decisions this framework supports. None of them change the framework's core recommendation — build the ecosystem before you scale autonomy — but they change the numbers, the regulatory clock, and the vendor landscape that boards will be asked to react to.

1. US adoption data got sharper — and it confirms shallow deployment

The US Census Bureau revised its Business Trends and Outlook Survey (BTOS) in November 2025 to measure AI use “in any business function” and fielded a second AI supplement (Nov 2025–Feb 2026). The results: 18–20% of US firms used AI in a business function through spring 2026, reaching 19.8% as of May 3, 2026, with 32% adoption on an employment-weighted basis and roughly 22% expecting to use AI within six months. Sector concentration is stark: Information (39.7%) and Finance and Insurance (33.9%) lead, while Retail Trade trails at ~14%. Most telling for this framework: among adopting firms, 57% use AI in three or fewer business functions. The majority of “AI adopters” have crossed the adoption line but not the integration line — which is exactly the gap between Level 2 and Level 3 on our maturity ladder.

2. Agentic AI reached the top of the hype curve

Gartner published its first dedicated Hype Cycle for Agentic AI in April 2026, placing the category at the Peak of Inflated Expectations. Its 2026 CIO survey found only 17% of organizations have actually deployed agents, against more than 60% planning to within two years — the most aggressive adoption-intent curve of any emerging technology it measures. Gartner also flagged “agent-washing” (rebranded automation sold as agents) as an explicit market problem and maintains that over 40% of agentic AI projects will be canceled by end-2027. Encouragingly, governance, security, and FinOps-for-agents now appear as named categories on the Hype Cycle itself — the market is learning, early, that the control plane is not optional.

3. The EU moved the high-risk deadline — but not the transparency deadline

On 7 May 2026, EU institutions reached political agreement on the Digital Omnibus on AI, the first amendments to the EU AI Act, with formal adoption completed in June 2026 (Parliament endorsement 16 June; Council approval 29 June). High-risk obligations for stand-alone Annex III systems (recruitment, credit scoring, education, law enforcement, and similar) are deferred from 2 August 2026 to 2 December 2027; obligations for AI embedded in regulated products (Annex I) move to 2 August 2028. Critically, the Article 50 transparency obligations — disclosing AI interactions to users — still apply from 2 August 2026, with a short grace period (to 2 December 2026) for machine-readable watermarking on systems already on the market. New Article 5 prohibitions on AI-generated non-consensual intimate imagery and CSAM take effect 2 December 2026, and SME compliance simplifications were extended to mid-caps (up to 750 employees / €150M revenue). Boards should read the deferral as breathing room to build compliance properly, not as permission to slow governance: the requirements did not change, only the enforcement date.

4. SAP's U-turn proved the transformation gap is real

At Sapphire 2026 (May, Orlando), SAP unveiled the “Autonomous Enterprise” and — notably — reversed its cloud-only AI posture, announcing that Joule assistants and agents will be made available to on-premises ECC and S/4HANA customers who commit at least 50% of their maintenance base to a cloud transition. This is a tacit acknowledgment of the numbers we cited in February: the majority of the ECC installed base will

not complete migration before the AI cycle matures, so the AI had to come to the legacy estate. The section on the Transformation Gap has been updated accordingly.

5. Interoperability standards matured

Protocol-level standardization accelerated: the Model Context Protocol (MCP) for agent-to-tool connectivity and agent-to-agent (A2A) interoperability schemes are now being adopted by major enterprise vendors, with Forrester predicting 30% of enterprise application vendors will ship MCP servers in 2026. For buyers, this modestly reduces platform-layer lock-in risk over time — and strengthens the case for the model-agnostic application layer this paper has always recommended.

The Executive Decision Framework: AI Adoption + Agent Autonomy

Decision you are making: Should we incorporate AI into the business? If yes, what level of autonomy and learning do we permit (copilots to agents to self-learning/self-modifying)?

Step	What You Decide	Key Takeaway
Step 1	Choose business outcomes (not AI projects)	If you cannot name measurable outcomes, you are not ready to adopt AI.
Step 2	Classify use cases by AI mode (Predict, Generate, Act)	Each mode has different risk and control requirements. Do not treat them the same.
Step 3	Select your maturity target (Level 0 through 5)	This is the real steering decision. Start low, graduate up only after building the ecosystem.
Step 4	Pass the ecosystem readiness gate	Each maturity level has minimum infrastructure you must have before moving up.
Step 5	Set autonomy and learning limits (“blast radius” contract)	Define what the AI can touch, how much it can do, and what is forbidden.
Step 6	Go/No-Go and funding rule	Fund in phases. Autonomy increases only after passing gates.

Step 1: Choose the Business Outcomes (not “AI projects”)

Pick 2 to 4 outcomes you will measure quarterly:

- Cost-to-serve reduction (cycle time, staffing load)
- Revenue lift (conversion, retention, upsell)
- Quality/error reduction (rework, write-offs, leakage)
- Risk reduction (fraud, compliance, incidents)
- Speed (SLA, time-to-decision)

If you cannot name measurable outcomes, stop here. You are not ready.

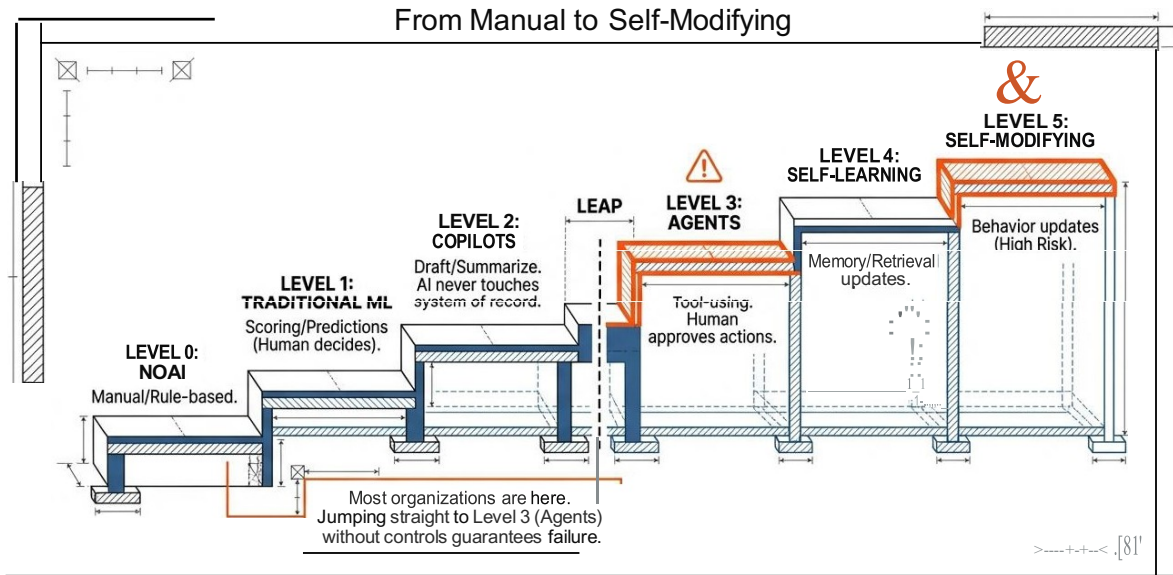
Step 2: Classify Candidate Use Cases by “AI Mode”

Put each use case into one of these buckets (because each has different risk and controls):

- **M1: Predict/Score (ML)** - forecasting, propensity, anomaly detection
- **M2: Understand/Generate (LLMs)** - summarization, drafting, Q&A, knowledge search
- **M3: Act (Agents)** - tool-using workflows that execute steps across systems

Step 3: Select the AI Maturity Target (12-18 months)

This is the real steering decision.



Level 0: No AI. Manual plus deterministic automation only.

This is where many organizations still are in production. All decisions and actions are performed by people or by traditional rule-based software (if X, then Y). There is no machine learning, no language models, and no intelligent automation. This is a perfectly valid starting point, and some processes should stay here permanently. Official business surveys that measure AI used in business functions (not just experimentation) show adoption is still a minority in several large economies — for example: ~20% of EU enterprises in 2025 (Eurostat), ~25% of UK businesses in late 2025 (ONS), and 19.8% of US businesses as of May 2026 (US Census Bureau BTOS; 32% employment-weighted, with ~22% expecting to use AI within six months). Adoption is higher for large firms and knowledge-intensive sectors — use rates reach 50–60% for very large firms in Information, Professional Services, and Finance — and definitions vary, so treat this as a baseline indicator, not a universal fact.

Level 1: Traditional ML at scale (M1). Scoring in batch/APIs; no autonomy.

At this level, the organization uses machine learning models to make predictions or flag patterns, but those predictions are delivered to humans who then decide what to do. For example, a model might score every customer by their likelihood to cancel their subscription, or flag transactions that look potentially fraudulent. The model generates a score or a recommendation; a person reviews it and acts. The AI has no ability to do anything on its own.

Level 2: LLM copilots (M2). Draft/summarize/retrieve with citations; no direct system writes.

Here, the organization introduces large language models that can read, write, and summarize text. A copilot might draft a response to a customer inquiry, summarize a 50-page contract, or search internal knowledge bases and return answers with citations. The critical constraint is that the copilot only produces text for a human to review. It cannot send an email, update a database, approve a transaction, or take any action in any system. A person always sits between the AI and the outside world.

Level 3: Tool-using agents with approvals (M3-A3). Agents can execute actions only after explicit approval; strict tool allowlists.

This is the first level where AI can do things in your systems, not just read or write text. An agent at this level might assemble a dispute response, look up the relevant policy, draft the letter, and then queue it for a human to review and approve before it gets sent. The agent has access to specific tools (like your CRM, ticketing system, or email), but every action requires explicit human approval before it executes. Think of it as a highly capable assistant that prepares everything but always asks before pressing “send.” As of 2026, this is where the frontier of responsible enterprise deployment sits: Gartner finds only 17% of organizations have deployed agents at all, and most production deployments remain narrowly scoped.

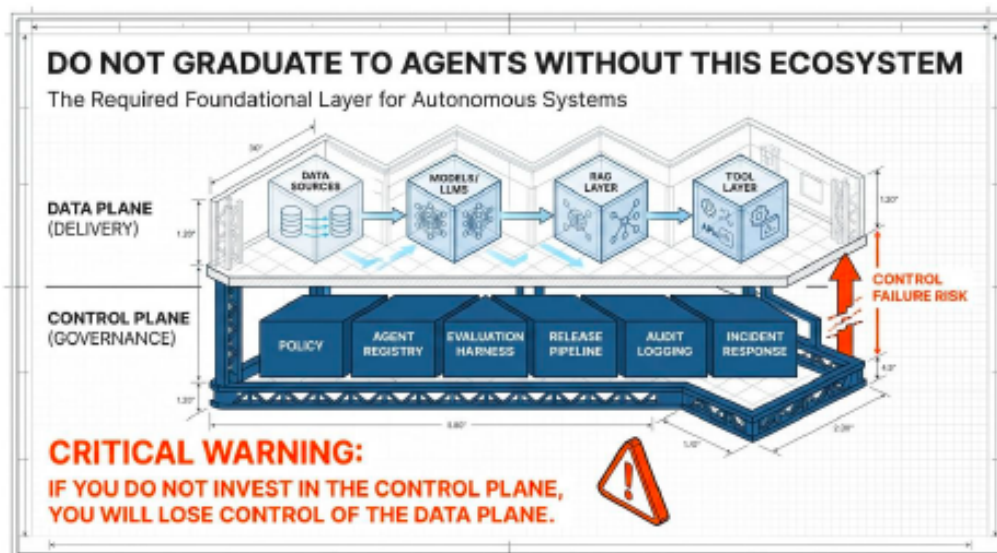
Level 4: Limited self-learning (Memory + Retrieval) (L1-L2). Agents can update scoped memory and retrieval indices; changes are reversible and audited.

At this level, the agent can learn from its interactions and improve over time within defined boundaries. For instance, it might update its own knowledge base when it encounters new information, refine how it searches for answers based on what worked before, or remember context from previous interactions with a specific customer. All these changes are logged, reversible, and scoped to specific areas. The agent is not rewriting its own instructions; it is updating the information it draws on.

Level 5: Controlled self-modification (Policies/Prompts) (L3). Agent behavior updates happen only through a gated release pipeline (tests, approvals, rollback). Weight updates and fine-tuning in production is usually a “no.”

This is the most advanced level, where the agent can modify how it behaves, not just what it knows. For example, it might update its own decision-making rules or adjust how it prioritizes tasks. This is powerful but carries the highest risk: an agent that can change its own behavior can drift in ways that are hard to detect. At this level, any change to agent behavior must go through the same kind of rigorous testing, approval, and rollback process that you would use for a production software release. No agent should be allowed to rewrite itself in real time without human oversight.

Step 4: AI Ecosystem Readiness Gate



Minimum ecosystem components you need before moving up:

For Level 1-2 (ML + Copilots):

- Data ownership and access controls
- Model/LLM usage policy and logging
- Basic evaluation and monitoring
- Security review for vendors/models

For Level 3 (Agents that act):

- Tool gateway, allowlisted actions, and schema validation
- Least privilege credentials and secrets management
- Audit logging (tamper-evident)
- Human approval workflow and separation of duties
- Circuit breakers / kill switch

For Level 4-5 (Learning/modifying):

- Versioned agent registry (charters, permissions, prompts/policies)
- Evaluation harness, regression, and adversarial testing
- Change control like software releases (two-person rule for high risk)
- Rollback and replay capability
- Incident response runbooks and on-call

The market is converging on this view: Gartner's 2026 Hype Cycle for Agentic AI now tracks agentic AI governance, agentic AI security, and FinOps for agentic AI as distinct technology categories — evidence that oversight infrastructure is being recognized as a precondition for scale, not an afterthought. Independent surveys underline why: substantial shares of organizations report they cannot enforce purpose limitations on agents, cannot quickly terminate a misbehaving agent, or cannot isolate AI systems from sensitive networks. Those are precisely the Level 3 gate items above.

Step 5: Set Autonomy and Learning Limits (the “blast radius” contract)

Autonomy: A0 read-only through A4 autonomous. Learning: L0 none through L4 model updates (offline only).

Why caps and approvals matter (risk math, not bureaucracy): if an agent has even a 0.01% per-action probability of a material policy breach or wrong write, then at 10,000 actions the chance of at least one incident is ~63% ($1 - (1 - 0.0001)^{10000}$). Action caps, approvals, and kill switches are the practical levers that keep N (actions) and p (per-action risk) bounded.

Define:

- Systems touched
- Action caps (per hour/day)
- Reversibility/rollback method

- Approval requirements
- Forbidden actions (hard deny)

Step 6: Go/No-Go and Funding Rule

- **Go if:** measurable ROI, controls in place, and operational ownership defined
- **No-Go if:** unmeasurable outcomes, missing telemetry, or you cannot bound blast radius
- Fund in phases. Autonomy increases only after passing gates.

The Business Case for “AI in General”

Where “AI in general” reliably pays off first

- Document and ticket summarization (support, ops, finance)
- Drafting and standard responses (customer comms, collections letters, internal updates)
- Knowledge retrieval with citations (policies, SOPs, product info)
- Prediction/scoring (churn risk, fraud signals, demand forecast)

Quantitative reality check: across multiple randomized and field experiments, copilots/LLM assistance can materially improve speed and quality, but results are heterogeneous. A widely cited field experiment in customer support found ~14% higher productivity on average, with larger gains among novice workers (Brynjolfsson et al., 2023). Controlled experiments on writing tasks report large time reductions (e.g., ~40% faster completion in one study), while other work shows a “jagged frontier” — performance improves on tasks within model strengths and can degrade outside them without training and scope boundaries. The 2026 US Census data adds a firm-level corroboration: firms with broader AI integration report stronger commercial performance and higher investment activity, but the majority of adopters remain narrow (three or fewer functions), and most use AI to augment tasks rather than replace headcount — AI-related employment decreases occurred in only ~2% of firms. This is why the framework insists on wedge use cases, workflow integration, and an evaluation discipline before scaling autonomy.

Common executive traps

- Treating AI as a tooling purchase rather than an operating model change
- Assuming copilots automatically produce ROI without workflow integration
- Ignoring evaluation/monitoring until after a high-profile failure
- Buying “agents” that are rebranded automation (“agent-washing”) and expecting agentic returns

What to approve first

Approve 1 to 2 “wedges” that:

- Reduce cycle time or cost-to-serve
- Fit existing workflows
- Have clean metrics
- Do not require high autonomy

Organizational Ownership and Accountability

The most common reason enterprise AI adoption stalls is not technical. It is organizational. In most companies, three or four executives each believe AI falls under their remit, and none of them want to be accountable if it goes sideways. The result is committee paralysis: endless pilot discussions, no production deployments, and a growing gap between the organization and its competitors.

Before selecting a single use case, leadership must answer a structural question: who owns AI in this organization? There are three viable models, and the right one depends on the company's size, culture, and operating model.

Model A: Centralized AI Office

A dedicated AI or Data Science team reports to a Chief AI Officer, Chief Data Officer, or CTO. This team owns the platform, sets standards, and partners with business units to deliver use cases. This model works well for organizations that need strong governance and are in regulated industries. The risk is that it becomes a bottleneck and disconnects from business priorities.

Model B: Federated with Central Standards

Each business unit has embedded AI capability (whether internal hires or vendor-supported), but a central team sets the architecture, security, and evaluation standards. The central team provides shared infrastructure (model hosting, tool gateways, evaluation harnesses) and reviews agents before production. This model works for larger organizations with diverse business lines. The risk is inconsistency and duplicated effort.

Model C: Business-Led with Technical Guardrails

A business leader (typically the COO or a P&L owner) owns AI outcomes, and technology provides the platform and guardrails. This model works when AI is being used primarily to optimize existing operations rather than build new products. The risk is that the business leader underestimates technical complexity.

Regardless of the model chosen, three accountability principles apply:

- A single executive must be named as the accountable owner for AI outcomes and risk. Shared ownership is no ownership.
- The accountable owner must have both budget authority and the ability to make production deployment decisions.
- Separation of duties must exist between those who build/configure agents and those who approve their deployment, particularly at Maturity Levels 3 and above.

Market practice is catching up to these principles: Forrester predicts 60% of Fortune 100 companies will appoint a dedicated head of AI governance in 2026. Whether that role sits inside the AI office or the risk function matters less than that it exists, is named, and has authority.

Executive Support: The Air Cover That Makes or Breaks AI Programs

Budget, resources, and support are the three prerequisites for any enterprise initiative to succeed. Most organizations get the first two partially right and the third one completely wrong. Budget gets allocated. Headcount gets approved. But the AI initiative owner is then left to navigate organizational resistance, skepticism, and the inevitable early failures alone. When the first pilot underperforms, when a model hallucinates something embarrassing, or when a business unit leader publicly questions the investment, the AI owner needs someone above them who will hold the line. Without that, the program dies not from technical failure but from political exhaustion.

Why Support Matters More Than Budget

AI adoption is not a linear process. Unlike a software implementation where you can project milestones with reasonable confidence, AI programs involve experimentation, iteration, and failure. The first model will not be good enough. The first agent will make mistakes. The first set of metrics will show ambiguous results. This is normal. It is how the technology works. But most organizations are not structured to tolerate this kind of ambiguity, especially when the investment is visible and the expectations are high.

The AI initiative owner needs the CEO, the board, or both to understand and communicate three things to the rest of the organization:

- AI success is non-linear. There will be a period of investment before returns materialize, and early results will be uneven. This is expected, not a sign of failure.
- The AI owner has the authority to make decisions, reallocate resources, and push back on business units that resist integration. This authority must be visible and reinforced, not just stated once in a memo.
- The organization will be measured on its willingness to adopt, not just the AI team's ability to deliver. If business units refuse to integrate AI tools into their workflows, that is a leadership problem, not a technology problem.

What CEO and Board Support Looks Like in Practice

Support is not a speech at a town hall. It is a set of repeated, concrete actions:

- Regular check-ins between the CEO and the AI owner (monthly at minimum) where progress is reviewed against the non-linear reality of the work, not against a waterfall timeline that was never realistic.
- Public reinforcement when setbacks occur. When a pilot underperforms, the CEO should frame it as a learning investment, not a failure. This sets the tone for the entire organization.
- Willingness to remove blockers. If a business unit leader is passively resisting adoption, the CEO must intervene directly. The AI owner should not have to fight political battles that only the CEO can win.
- Board-level reporting that includes AI progress as a standing agenda item, not a one-time update. This signals to the entire C-suite that AI is a strategic priority, not a side experiment.
- Protection of the learning period. For the first 6 to 12 months, the AI program should be evaluated on capability-building milestones (models deployed, workflows integrated, evaluation infrastructure built)

rather than hard ROI targets. Premature ROI pressure kills AI programs before they have a chance to compound.

The Consequences of Absent Support

When the AI owner lacks genuine executive support, a predictable pattern emerges. Early wins are achieved through sheer effort, but they are fragile. The first visible failure triggers skepticism. Budget gets quietly redirected. The best technical talent leaves because they can see the program is not backed. Within 12 to 18 months, the organization declares that “AI did not work for us” when what happened is that leadership did not create the conditions for it to work. Gartner's projection that more than 40% of agentic AI projects will be canceled by end-2027 is, in our experience, as much a forecast of leadership failure as of technology failure.

The enterprises that succeed with AI are not the ones with the biggest budgets or the most advanced technology. They are the ones where the person running the AI program has a CEO who says, publicly and repeatedly, “This is a priority, this person has my backing, and we are going to stay the course even when it gets uncomfortable.” That is what support means.

Talent and Operating Model

The team required to run AI well changes significantly as you move up the maturity ladder. Most enterprises underestimate this. They fund the technology but not the people who keep it working. Below is a practical guide to the roles required at each level.

Levels 1-2 (ML and Copilots)

At this stage, you need data engineers who can build and maintain clean data pipelines, a machine learning engineer or applied scientist who can evaluate and fine-tune models, a product or program manager who can translate business outcomes into use case specifications, and a security/compliance reviewer who can assess vendor data handling. Many organizations can start here with 3 to 5 people plus vendor support.

Level 3 (Agents That Act)

Agents require a meaningful step up in operational capability. In addition to the Level 1-2 team, you need platform/infrastructure engineers who can build and maintain tool gateways, approval workflows, and audit logging. You need an evaluation specialist who can design test harnesses and adversarial scenarios. You also need an on-call rotation with incident response capability, because agents can fail at any hour. Note that the skills gap here is now a named market constraint: Gartner identifies eval-driven agent development as a scarce blend of software and AI engineering skills, and warns that building customized agent stacks with pro-code frameworks can produce overly complex systems that are difficult to support long-term.

Levels 4-5 (Learning and Self-Modification)

At this level, the team resembles a software release engineering function. You need release engineers who manage the gated deployment pipeline for agent behavior changes. You need adversarial testers who probe for prompt injection, tool abuse, and drift. And you need a governance lead who maintains the agent registry, reviews charters, and owns the audit trail.

Build vs. Buy vs. Partner

Most enterprises at Levels 1-3 will rely heavily on vendor platforms and system integrators. The critical question is not “build or buy” but “what must we own internally.” At a minimum, you must own the evaluation and monitoring capability. If you outsource the ability to know whether your AI is working correctly, you have outsourced the ability to manage risk.

When (and Why) to Move from Copilots to Agents

What changes when you introduce agents

Agents do not just generate text. They trigger actions across systems. That requires:

- Access control
- Approval gates
- Audit logs
- Kill switches
- Deterministic tool interfaces

The 2026 reality check

The gap between agent ambition and agent deployment is now well documented. Gartner's 2026 CIO survey found only 17% of organizations have deployed AI agents, while more than 60% plan to within two years. Industry surveys consistently show adoption claims far outrunning production reality — one widely cited analysis found 79% of enterprises claiming to have “adopted AI agents” while only around one in ten run them in production. Much of what is marketed as agentic is “agent-washed” automation. Two implications for boards: first, competitor press releases about agents are a poor benchmark for your own urgency; second, when your organization does move, the constraint will be the control plane (approvals, audit, evaluation, incident response), not the model. Interoperability standards such as the Model Context Protocol (MCP) are lowering the technical cost of connecting agents to enterprise systems — which makes disciplined tool allowlisting and gateways more important, not less, because it becomes easier for an over-permissioned agent to reach systems it should not.

Good first agent targets (bounded, reversible)

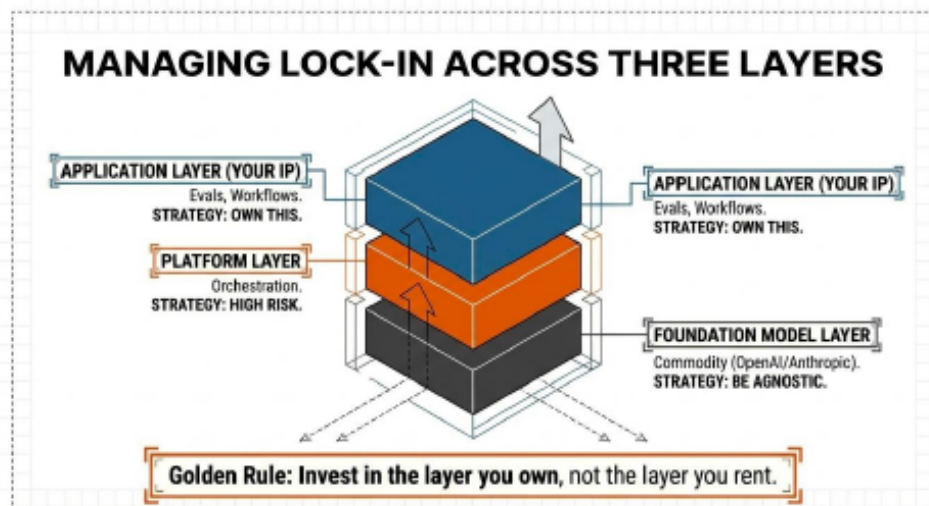
- Service desk runbook execution (approval required)
- Disputes/collections case assembly (agent drafts, human approves sends/updates)
- Procurement intake (agent classifies, drafts comparisons, routes approvals)
- Data operations (agent proposes fixes, generates scripts, but cannot deploy)

Where you should not start

- Autonomous customer communications without review
- Anything that can move money, change pricing, approve credit, or alter financial postings without multi-step approvals.

The Vendor Decision

For most enterprises at Maturity Levels 1 through 3, the AI adoption decision is, in practice, a vendor selection decision. You are not building foundation models. You are choosing which ecosystem to commit to, how much lock-in you are willing to accept, and where you draw the line between platform dependency and internal capability.



The Three Layers of Vendor Dependency

Think about vendor relationships across three layers:

Foundation model layer. This is where you choose between providers like OpenAI, Anthropic, Google, Meta (open-source), or others. The key risk is that model capabilities, pricing, and data handling policies change frequently. The mitigation is to design your application layer to be model-agnostic where possible, so you can swap providers without rebuilding workflows.

Platform/orchestration layer. This includes the tools that manage prompts, retrieval, agent orchestration, and tool connectivity. Vendor options range from cloud-native solutions (Azure AI, AWS Bedrock, Google Vertex) to independent platforms. Lock-in risk is highest here because switching platforms means re-engineering workflows. Two 2026 developments modestly soften this risk over time: the spread of the Model Context Protocol (MCP) as a standard way to connect agents to data and tools, and emerging agent-to-agent (A2A) interoperability between major vendors' agent frameworks. Ask every platform vendor about their support for open protocols; it is now a legitimate selection criterion.

Application layer. This is where AI meets your specific business processes. It includes the integrations with your systems of record, the approval workflows, the evaluation harnesses, and the monitoring dashboards. This layer should be owned internally whenever possible, even if you use vendor components underneath.

Non-Negotiable Vendor Evaluation Criteria

- **Data usage and retention:** Will the vendor train on your data? Can you opt out? What is the retention policy?
- **Audit export:** Can you export complete logs of all model interactions, tool calls, and agent actions?
- **Evaluation tooling:** Does the vendor support systematic evaluation, or are you limited to manual spot-checking?
- **Cost transparency:** Is pricing predictable? Are there usage-based spikes that could surprise finance?
- **Exit strategy:** What does migration look like if you need to switch providers? Does the vendor support open interoperability standards (e.g., MCP)?
- **Agent substance:** If the vendor sells “agents,” can they demonstrate autonomous planning, tool orchestration, and persistent context — or is it scripted automation with a new label?

The Required AI Ecosystem (What You Must Build)

Think of the ecosystem as two planes:

1) Data-plane (delivery)

- Data sources, ingestion, and access controls
- Models/LLMs and prompt/policy layer
- Retrieval layer (RAG)
- Agent/tool layer (for agents)
- Monitoring and telemetry

2) Control-plane (governance)

- AI policy and risk management
- Model/agent registry (versions, permissions, charters)
- Evaluation harness (golden tests and adversarial tests)
- Release pipeline and rollback
- Incident response and audit exports

If you do not invest in the control-plane, you will eventually lose control of the data-plane.

Risk, Compliance, and Non-Negotiable Controls

AI risks (general)

- Data leakage, IP exposure, privacy violations
- Hallucinations and incorrect recommendations
- Biased or inconsistent outcomes
- Vendor lock-in and uncontrolled cost

Agent risks (additional)

- Prompt injection through docs/tickets
- Tool abuse via over-permissioned tokens
- Silent drift (learning makes it worse over time)
- Reputational damage at scale

Note: these are not theoretical risks. OWASP's guidance for LLM applications and AI agents explicitly flags prompt injection, insecure output handling, and excessive agency/tool access as top failure modes.

Separately, primary research has demonstrated indirect prompt injection: malicious instructions embedded in documents/web pages can be retrieved into an agent's context and executed as if they were trusted instructions. If you plan to index tickets, emails, or documents for RAG, you must treat that content as untrusted input and test accordingly. Perception has caught up with the research: a 2026 Gartner survey of IT application leaders found 74% view AI agents as a new attack vector — yet only 13% strongly agreed they had the right governance structures in place. That 61-point gap is the single clearest quantification of the control-plane deficit this paper exists to close.

Regulatory update (Q3 2026): the EU AI Act after the Digital Omnibus

The most significant regulatory development of 2026 is the Digital Omnibus on AI — the first amendments to the EU AI Act — politically agreed on 7 May 2026 and formally adopted in June 2026. What boards need to know:

- **High-risk deadlines deferred.** Obligations for stand-alone Annex III high-risk systems (recruitment, credit scoring, education, law enforcement, and similar use cases) move from 2 August 2026 to 2 December 2027. Obligations for AI embedded in regulated products (Annex I — medical devices, machinery, vehicles) move to 2 August 2028.
- **Transparency obligations were NOT deferred.** Article 50 obligations — disclosing to users that they are interacting with an AI system — apply from 2 August 2026. Machine-readable watermarking of AI-generated content gets only a short grace period, to 2 December 2026, for systems already on the market. Fines for violations can reach €15 million or 3% of worldwide turnover.
- **New prohibitions.** AI systems that generate or manipulate non-consensual intimate imagery or child sexual abuse material are prohibited effective 2 December 2026, added to the Article 5 prohibited-practices list (violations of prohibitions carry fines up to €35 million or 7% of turnover).
- **Simplification for mid-caps.** SME compliance simplifications (documentation templates, proportionate penalties, sandbox access) were extended to companies with up to 750 employees and €150 million revenue.

- **Scope refinement.** AI components that merely assist users or optimize performance — without creating health or safety risk on failure — will not qualify as high-risk “safety components,” narrowing over-inclusive classification.

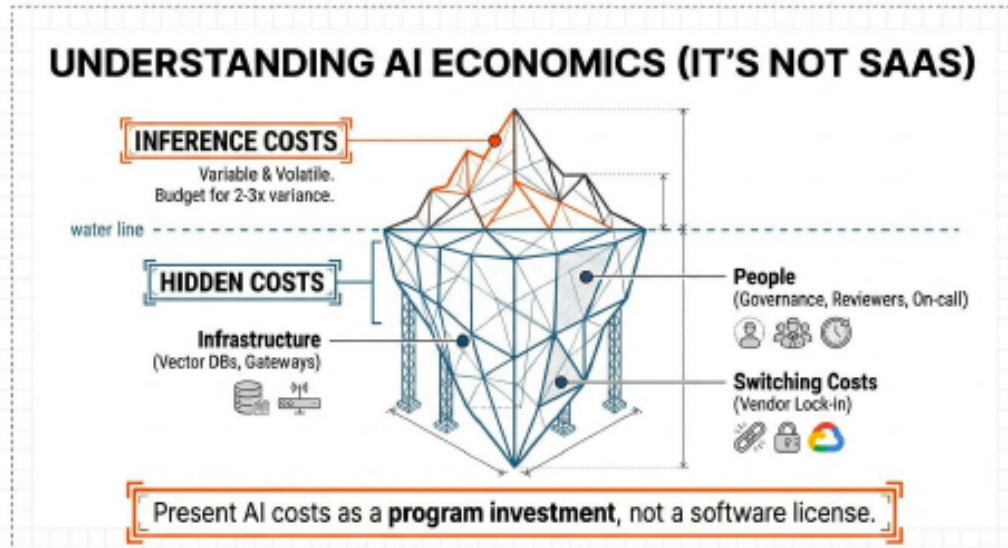
Our guidance: treat the deferral as time to do the work properly, not as a reason to deprioritize it. The substantive requirements — risk management, testing, logging, human oversight, post-market monitoring — did not change, and they map almost one-to-one onto the ecosystem readiness gates in this framework. An organization that builds the Level 3 control plane described in this paper will find EU AI Act high-risk compliance largely a documentation exercise. An organization that waits until mid-2027 to start will be rebuilding its agent architecture under deadline pressure.

Non-negotiables to require before “agents that act”

- Least privilege
- Tool allowlists and strict schemas
- Audit logs with tamper evidence
- Human approvals for high-impact actions
- Circuit breakers and kill switch
- Release discipline (no ad hoc “agent changes”)

Understanding the AI Cost Model

AI costs do not behave like traditional enterprise software licensing, and this catches many executives off guard. A SaaS platform has a predictable per-seat or per-module cost. AI costs are usage-driven, variable, and often opaque until you are already committed.



The Four Cost Categories

1. Inference costs. Every time a model processes a request, you pay for tokens (the unit of text processed). Costs vary dramatically by model: a simple summarization call on a smaller model may cost fractions of a cent, while a complex multi-step agent workflow using a frontier model can cost dollars per execution. At enterprise scale, these costs compound quickly. Budget for usage-based pricing with a meaningful variance range (often 2-3x your initial estimate during ramp-up). Public cloud pricing pages (e.g., AWS Bedrock) make this explicit: pricing is typically per-token/per-request and can change as models and tiers evolve. Note that agentic workloads change the cost shape: an agent that plans, retrieves, calls tools, and self-checks can consume an order of magnitude more tokens per task than a single copilot prompt — which is why FinOps for agentic AI now appears as its own category on Gartner's 2026 Hype Cycle.

2. Infrastructure and platform costs. These include the hosting environment for your retrieval layer (vector databases, search indices), the tool gateway and orchestration layer, logging and monitoring infrastructure, and evaluation/testing environments. These costs are more predictable but are frequently underestimated because they are treated as "technical overhead" rather than core program costs.

3. People costs. The talent required to run AI well is the largest ongoing cost at Maturity Levels 3 and above. This includes not just the build team but the operational team: evaluation specialists, on-call engineers, governance leads, and the business analysts who define and measure outcomes. Most failed AI programs are underfunded on people, not technology.

4. Opportunity and switching costs. Choosing a vendor ecosystem creates dependency. Switching costs are real and include re-engineering integrations, retraining staff, re-validating outputs, and potential data migration. Factor these into your total cost of ownership from the start, not as an afterthought.

CFO-Ready Guidance

When building the business case, present AI costs as a program investment (analogous to a digital transformation initiative), not a line-item software purchase. Include a 6-month ramp-up period where costs will be higher relative to output as the team builds capability. Set quarterly cost review gates tied to the same KPIs used to measure business outcomes. If the cost-to-value ratio is not improving quarter over quarter, that is an early warning signal that the program needs course correction.

Roadmap and Funding Logic (12-18 months)

Phase 0 (2-4 weeks): Decide and set guardrails

- Pick outcomes and KPIs
- Choose 1-2 wedges (M1 or M2)
- Stand up AI governance and logging
- Define autonomy/learning limits (future-proof)

Phase 1 (4-10 weeks): Deliver measurable wins (copilots/ML)

- Deploy copilots in workflows (draft and retrieve with citations)
- Establish baseline vs. post metrics
- Build evaluation harness

Phase 2 (8-16 weeks): Introduce bounded agents (A3)

- Tool gateway, approvals, audit, and kill switch
- Start with reversible actions and strict caps

Phase 3 (ongoing): Enable limited learning (L1-L2)

- Memory and retrieval updates only
- Full provenance, TTL, and rollback

Phase 4 (rare): Controlled self-modification (L3)

- Gated prompt/policy updates only through CI/testing and approvals
- No “live self-rewrite” in production.

Regulatory checkpoint for EU-exposed organizations: map every deployed system against the amended EU AI Act calendar — transparency obligations from 2 August 2026, watermarking from 2 December 2026, Annex III high-risk from 2 December 2027, Annex I from 2 August 2028 — and align your Phase 2/3 control-plane build so that conformity evidence (logs, testing, oversight records) is produced as a by-product of operations, not a separate project.

The Transformation Gap: What Enterprises Actually Face

Enterprises are not two years away from being AI-native. When the full transformation sequence is mapped honestly, most large organizations are four to six years away. The gap between those two estimates is almost entirely organizational, not technical. Understanding why requires mapping the actual sequence of work, not the marketing sequence. Insight: The delay is not caused by new technology; it is caused by the pace at which organizations can absorb new decision-making technology.

Start with the infrastructure problem. Gartner data shows that at the end of 2024, only about 39% of the roughly 35,000 SAP ECC customers worldwide had migrated to S/4HANA — the platform that is a prerequisite for the full AI capabilities SAP markets — and independent analyses project that fewer than 57% will complete migration before ECC's 2027 end-of-mainstream-support deadline. A 2026 survey of German-speaking SAP user groups (DSAG) found 78% of members operating hybrid landscapes, 54% still running ECC or the old Business Suite somewhere in their estate, and nearly half of remaining ECC users planning migration by 2030, not 2027. You cannot apply an AI layer to a system you have not migrated. The migration itself is a two-to-four-year program — Gartner has worked with clients on three- to seven-year projects — costing anywhere from \$2 million to, for the most complex estates, up to \$1 billion before a single AI feature is configured or tested.

The 2026 plot twist proves the point. At Sapphire 2026 (May, Orlando), SAP unveiled the “Autonomous Enterprise” — a unified Business AI Platform, an autonomous suite, and Joule agents built with foundation-model partners including Anthropic — and simultaneously reversed its long-held cloud-only AI posture. Joule assistants and agents will now be made available to on-premises ECC and S/4HANA customers, but only for those who commit at least 50% of their maintenance base to a cloud transition. Read plainly: the vendor with the most complete view of enterprise back-office reality concluded that its installed base will not finish migrating before the AI cycle matures, so the AI must be brought to the legacy estate as a bridge. SAP also introduced agent-led migration tooling it claims cuts ERP migration effort by more than 35% — AI being deployed to accelerate the very migrations that AI adoption was waiting on. Both moves validate the transformation-gap thesis: availability was never the constraint; absorption was.

The headline numbers still compound the confusion. SAP reports Joule embedded in over 80% of its most-used tasks, with Joule adoption growing ninefold during 2025 and more than 2,400 prebuilt skills across its applications. But “embedded” and even “adopted ninefold” mean available and activated in the product — not deployed, configured, tested, validated against a specific enterprise's data architecture and compliance requirements, and actually absorbed by a workforce whose incentive structures and decision-making habits were built for a pre-AI world. Available and adopted are separated by the entirety of the hard work. The same pattern shows up economy-wide in the 2026 US Census data: among firms that report using AI at all, 57% use it in three or fewer business functions.

The actual sequence a large enterprise faces is rarely articulated plainly, but it runs as follows. First, the enterprise must be on a platform version that supports the intended AI capabilities — or negotiate bridge access to AI on its legacy estate, as SAP now offers. Most are not there. Second, they must complete platform migration, typically two to four years depending on system complexity and technical debt. Third, they must configure AI features for their specific workflows, data architecture, and regulatory environment. Fourth, they must train the workforce, not just on the tool but on new ways of working. Fifth, and this is

where most roadmaps simply stop, they must actually change how decisions get made. That last step is not a technology project. It is a leadership project, and it takes longer than any software upgrade.

Enterprises that attempt to skip from foundational automation directly to agentic or self-modifying systems, without first proving reliability, auditability, and organizational trust at lower capability levels, tend to roll back the entire program. Gartner's standing prediction that more than 40% of agentic AI projects will be canceled by end-2027 is the market-level expression of this pattern. The sequencing is not a theoretical construct. It is a practical constraint imposed by the pace at which organizations can absorb and govern new decision-making infrastructure. Any honest autonomy and displacement strategy must force leadership to make explicit choices about how much autonomy the organization is ready to control, not how much the technology could theoretically support.

A company can have AI embedded in its ERP environment, hold licenses for every major AI tool across its software stack, and still not be AI-native, because its processes, incentive structures, and governance culture were designed for a world in which humans made every consequential decision and software simply recorded the outcomes. Redesigning that operating model is not a function of software procurement. It is a function of organizational transformation at the executive level, which proceeds at a fundamentally different pace than technology deployment.

The enterprises most likely to close the gap are those whose leadership teams understand this distinction early, build the organizational infrastructure before scaling autonomy, and resist the pressure to conflate tool availability with operational readiness. The ones most likely to lose years are those whose AI strategy is owned by IT rather than by the executive accountable for the business outcome it is meant to change.

Executive Appendix: Toolkit

The following templates are designed to accelerate your AI adoption program. Each one has been refined through real enterprise engagements and maps directly to the decision framework outlined in this paper.

- **AI Strategy One-Pager:** outcomes, KPIs, wedge use cases, maturity target, budget bands
- **Use Case Scorecard:** value, feasibility, control feasibility, observability, time-to-impact, risk tier
- **AI Ecosystem Readiness Checklist:** data governance, logging, evals, monitoring, IAM, vendor posture
- **Copilot Deployment Checklist:** grounding rules, citations, redaction, user training, human QA loop
- **Agent Charter Template:** purpose, autonomy tier, learning level, allowed tools/actions, forbidden actions, approvals, caps, rollback
- **Action Allowlist/Denylist Template:** tool endpoints, objects, fields, rate limits, approval requirement
- **Risk Register Template:** top 10 risks with detection and response
- **Pilot Exit Criteria:** KPI thresholds, safety thresholds, and operational readiness
- **Vendor Evaluation Checklist:** data usage/retention, training on your data, audit export, controls, eval tooling, cost transparency, open-protocol support
- **Organizational Ownership Matrix:** RACI chart for AI decisions across CTO, COO, business unit leads, and compliance
- **AI Cost Model Template:** inference, infrastructure, people, and switching cost projections by maturity level
- **Talent Requirements by Maturity Level:** role descriptions, headcount ranges, and build-vs-buy guidance
- **EU AI Act Compliance Calendar (new):** system inventory mapped against the amended deadlines — transparency (Aug 2026), watermarking (Dec 2026), Annex III high-risk (Dec 2027), Annex I (Aug 2028)

To receive the complete toolkit with ready-to-use templates, or to discuss how this framework applies to your organization, contact us at tk@5054holdings.com or Cyril.simone@dimensioneleven.com. We work with enterprises at every stage of the AI maturity ladder, from first-use-case selection through full agent deployment.

About the Authors

Tunde Kehinde is the Managing Partner of 5054 Holdings, where he invests in and builds businesses across technology and infrastructure. He co-founded Jumia, helping scale the business through its formative growth years and contributing to over \$1 billion in enterprise value by designing go-to-market strategy, operating cadence, and trust mechanisms in complex, low-trust markets prior to its public listing on the NYSE. He subsequently led growth and investment strategy at Lidya, a cross-border fintech lender, where he originated and deployed over \$150 million in private credit across the U.S., Europe, and Africa, achieving a 99% repayment rate through rigorous underwriting, dynamic pricing, and disciplined borrower acquisition. Across his career, Tunde has led over 1,500 direct and indirect reports across multiple continents, installing the operating rhythm, ownership structures, and performance standards required to sustain growth at scale. He is a Guest Lecturer at Howard University School of Business and serves as an advisor to enterprises navigating AI adoption, operating model transformation, and capital allocation strategy. Tunde completed his MBA at Harvard Business School.

Cyril Simone is a technology executive and researcher with over 20 years of experience at the intersection of AI, machine learning, and enterprise transformation. He served as VP of Global Technology at News Corp, where he led large-scale technology strategy across the company's media and publishing portfolio. Prior to that, he was Head of Northeast Professional Services at Sun Microsystems and served as Chief Architect for a multi-year transformative education product for the K-12 space in the United States, a joint venture with Office Depot and IBM Watson. As Senior Manager of Data Science at Capgemini, Cyril was Sector Lead for media and entertainment clients including Comcast/NBCUniversal, Disney, Sony, CBS, and Viacom. He has managed global teams of over 500 people. Cyril currently leads research in non-linear and quantum computing solutions for AI performance optimization, and is the founder of iLiner Studios, which creates immersive digital music experiences using advanced AI, and Mercado AI, an AI-driven marketplace that optimizes the agricultural supply chain from farm to market.

Bibliography

The following sources are cited in this paper. All URLs verified as of July 2026.

Governance frameworks, standards, and regulation

NIST. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1.
<https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1.
<https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>

ISO/IEC 42001:2023. Artificial intelligence management system (AIMS) requirements.
<https://www.iso.org/standard/42001.html>

European Union. Artificial Intelligence Act (EU AI Act), Official Journal publication (2024/1689). https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ%3AL_202401689

Covington & Burling (Inside Global Tech). “EU AI Act Update: Timeline Relief, Targeted Simplification, and New Prohibitions.” May 2026. Source for the Digital Omnibus on AI provisional agreement (7 May 2026), the deferral of Annex III high-risk obligations to 2 December 2027 and Annex I to 2 August 2028, and the new Article 5 prohibitions effective 2 December 2026. <https://www.insideglobaltech.com/2026/05/28/eu-ai-act-update-timeline-relief-targeted-simplification-and-new-prohibitions/>

Gibson Dunn. “EU AI Act Omnibus Agreement — Postponed High-Risk Deadlines and Other Key Changes.” May 2026.
<https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/>

Latham & Watkins. “AI Act Update: EU Resolves to Change Rules and Extend Deadlines.” May 2026. Source for Article 50 transparency obligations applying from 2 August 2026, the watermarking grace period to 2 December 2026, and associated penalty levels. <https://www.lw.com/en/insights/ai-act-update-eu-resolves-to-change-rules-and-extend-deadlines>

DLA Piper (GENIE). “The Digital AI Omnibus: Proposed Deferral of High-Risk AI Obligations under the AI Act.” Updated June/July 2026. Source for European Parliament endorsement (16 June 2026) and Council final approval (29 June 2026). <https://knowledge.dlapiper.com/dlapiperknowledge/globalemploymentlatestdevelopments/2026/The-Digital-AI-Omnibus-Proposed-deferral-of-high-risk-AI-obligations-under-the-AI-Act>

Empirical productivity evidence (selected causal studies)

Brynjolfsson, Li, and Raymond (2023). Generative AI at Work. NBER Working Paper 31161.
<https://www.nber.org/papers/w31161>

Noy and Zhang (2023). Experimental evidence on the productivity effects of generative artificial intelligence. Science.
<https://www.science.org/doi/10.1126/science.adh2586>

Dell'Acqua et al. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. HBS working paper.
https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf

Peng et al. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot (controlled experiment). arXiv:2302.06590. <https://arxiv.org/abs/2302.06590>

Brynjolfsson, Rock, and Syverson (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. AEA: Microeconomics. <https://www.aeaweb.org/articles?id=10.1257/mac.20180386>

Bubeck et al. (2026). Executive survey evidence on AI adoption and measured productivity impacts (cross-country). NBER Working Paper 34836. https://www.nber.org/system/files/working_papers/w34836/w34836.pdf

Agentic AI market and deployment evidence

Gartner. 2026 Hype Cycle for Agentic AI. April 2026. Source for agentic AI at the Peak of Inflated Expectations; the 2026 CIO and Technology Executive Survey finding that 17% of organizations have deployed AI agents while 60%+ expect to within two years; the emergence of agentic AI governance, security, and FinOps categories; and the “agent-washing” warning. <https://www.gartner.com/en/articles/hype-cycle-for-agentic-ai>

Gartner. “Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025.” Press release, August 2025. <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>

Security guidance for LLMs and agents

OWASP. Top 10 for Large Language Model Applications. <https://owasp.org/www-project-top-10-for-large-language-model-applications/>

OWASP Cheat Sheet Series. AI Agent Security Cheat Sheet. https://cheatsheetseries.owasp.org/cheatsheets/AI_Agent_Security_Cheat_Sheet.html

Primary research on indirect prompt injection against LLM-integrated applications. arXiv:2302.12173. <https://arxiv.org/abs/2302.12173>

Official enterprise AI adoption statistics (definition-dependent)

US Census Bureau. “Large Firms With at Least 20 Employees Biggest AI Users.” America Counts, May 2026. Source for BTOS AI-use rates of 17–20% (Dec 2025–May 2026), 19.8% as of May 3, 2026, and sector rates for Information (39.7%), Finance and Insurance (33.9%), and Retail Trade (~14%). <https://www.census.gov/library/stories/2026/05/ai-use-businesses.html>

US Census Bureau. “The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks.” Working Paper CES-WP-26-25, April 2026. Source for 18% firm adoption (32% employment-weighted) during the Nov 2025–Jan 2026 supplement period, 57% of adopters using AI in three or fewer functions, 50–60% use rates among very large firms in knowledge-intensive sectors, ~66% augmentation-only use, and AI-related employment decreases in only ~2% of firms. <https://www.census.gov/library/working-papers/2026/adrm/CES-WP-26-25.html>

Federal Reserve Board. “Monitoring AI Adoption in the U.S. Economy.” FEDS Notes, April 2026. <https://www.federalreserve.gov/econres/notes/feds-notes/monitoring-ai-adoption-in-the-u-s-economy-20260403.html>

Eurostat (2025). Share of enterprises using artificial intelligence technologies (EU). <https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251211-2>

UK Office for National Statistics (ONS) (Jan 2026 bulletin reporting late-2025 results). Business insights and impact on the UK economy. <https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/bulletins/businessinsightsandimpactontheukeconomy/8january2026>

Cost and pricing reference (example)

AWS. Amazon Bedrock pricing (example of per-token/per-request pricing model). <https://aws.amazon.com/bedrock/pricing/>

SAP ECC to S/4HANA migration and Business AI

Gartner (via CIO.com). “Nearly Half of SAP ECC Customers May Stick with Legacy ERP Beyond 2027.” CIO, 2025. Source for the finding that at the end of 2024 only ~39% of ~35,000 ECC customers had migrated to S/4HANA, and for migration cost (\$2M to \$1B) and duration (three- to seven-year) ranges.

<https://www.cio.com/article/4000543/nearly-half-of-sap-ecc-customers-may-stick-with-legacy-erp-beyond-2027.html>

Gartner (via The Register). "Legacy SAP ERP Customers Still Foot-Dragging Move to S/4HANA." The Register, March 2025. https://www.theregister.com/2025/03/20/sap_sees_little_progress_in/

Basis Technologies. "The True State of S/4HANA 2025." Independent analysis projecting fewer than 57% of ECC customers will complete migration by the 2027 end-of-mainstream-support deadline. <https://www.basistechnologies.com/blog/the-true-state-of-s4hana/>

SAP News Center. "SAP Unveils the Autonomous Enterprise." SAP SE, May 2026 (Sapphire Orlando). Source for the Business AI Platform / Autonomous Suite launch, foundation-model partnerships (including Anthropic), extension of Joule assistants and agents to on-premises ECC/S/4HANA customers committing the majority of their landscape to cloud, and agent-led migration tooling claiming 35%+ effort reduction. <https://news.sap.com/2026/05/sap-sapphire-sap-unveils-autonomous-enterprise/>

The Register. "SAP U-turn Brings AI Features to ECC and On-Prem S/4HANA." May 2026. Source for the 50%-of-maintenance-base cloud-commitment condition and Max Success Plan packaging. <https://www.theregister.com/ai-ml/2026/05/13/sap-u-turn-brings-ai-features-to-ecc-and-on-prem-s4hana/5239040>

SAP News Center. "SAP Announces Q4 and FY 2025 Results." SAP SE, January 2026. Source for ninefold Joule adoption growth during 2025 and 2,400+ prebuilt skills. <https://news.sap.com/2026/01/sap-announces-q4-and-fy-2025-results/>

SAP News Center. "SAP Sapphire Orlando Keynote: Bring Out Your Best." SAP SE, May 2025. Source for the unveiling of SAP AI Foundation at Sapphire 2025 and the statement that 80% of SAP's most-used transactions are managed via Joule. <https://news.sap.com/2025/05/sap-sapphire-companies-facing-big-challenges-we-are-on-your-side/>

IgniteSAP. "SAP Financial Results Q1 2026." April 2026. Source for DSAG survey findings: 78% hybrid landscapes, 54% still running ECC/Business Suite, and nearly half of remaining ECC users planning migration by 2030. <https://ignitesap.com/sap-financial-results-q1-2026/>