

AI Strategy Landscape

From Adoption to Advantage

Gerald C. Kane, Ph.D.

Executive MBA course, Terry College of Business, University of Georgia
Fall 2026

This public edition omits the in-class activity specifications, which are designed to be experienced live in session.

© 2026 Gerald C. Kane. Licensed under CC BY-NC 4.0. You may share and adapt it for noncommercial use with credit to the author; you may not present it as your own. Full terms on the last page.

Module role: Strategic framing module. The module establishes the 2026 AI business landscape and gives executives a shared vocabulary for discussing AI advantage, value, and competition. Subsequent modules go deeper on executive AI literacy, use-case discovery, stack and sourcing, governance, economics, capability acceleration, and the capstone action plan.

Module Learning Objective

By the end of the module, you will be able to explain the 2026 AI business landscape, distinguish capability adoption from enterprise value creation, and identify where durable competitive advantage may exist — through proprietary context, workflow integration, organizational learning, customer relationships, and AI-native business-model design.

If You Want to Know More

The Pre-Class Read below carries the argument; the materials below carry the underlying data and provocation. Read what fits your time:

- Stanford HAI, *2026 AI Index Report* — economy and technical-performance highlights.
 - McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation*.
 - MIT Sloan Management Review, either "Why AI Will Not Provide Sustainable Competitive Advantage" or "Generate Value From GenAI With 'Small t' Transformations."
 - One current business example you find compelling — e.g., Amazon Rufus, Colgate-Palmolive's GenAI program, Harvey, Cursor, or Perplexity.
-

Part 1: Pre-Class Read

BOTTOM LINE. AI capability is becoming broadly accessible — Stanford's 2026 AI Index puts organizational adoption at 88% and generative AI use in at least one business function at 70%, and McKinsey's 2025 survey reports a similar pattern. Access, however, is not advantage. Most organizations report little or no enterprise-level financial impact yet. The strategic question for 2026 is no longer whether to adopt AI; it is how to convert AI usage into durable enterprise value. Advantage tends to accrue to organizations that pair AI capability with proprietary context, workflow integration, organizational learning loops, and earned trust — not to those that simply have the best model.

Opening Vignette: Aflac and the Architecture of Discipline

Aflac, the supplemental insurance company headquartered in Columbus, Georgia, processes claims at a scale and pace where automation has obvious appeal — and where mishandling a claim is paid in customer trust during a difficult moment in a policyholder's life. The company built one of the earliest large-scale AI claims platforms in its industry, deployed in 2022 and 2023, and has since extended automation across accident, hospital, cancer, dental, and vision lines.

Shelia Anderson, then chief information officer at Aflac, described the design choice on MIT Sloan Management Review's *Me, Myself, and AI* podcast: "We had about 46 of our claims that were fully automated, or what we call straight-through processing — primarily wellness claims... 49% of our manual claims result in payouts less than, for example, \$200" (*Shelia Anderson, CIO, Aflac, MIT Sloan Management Review podcast, 2023*). The strategic logic behind the number is the part that matters. Aflac targeted automation at a specific layer — repetitive low-value claims where speed of resolution matters and the risk of error is bounded — and explicitly preserved human judgment for the higher-stakes complex cases where empathy, edge-case discretion, and accountability shape the experience. By 2025, more than 55% of claims were running through some form of AI assistance (*Keith Farley, Vice President of Innovation Operations, Aflac, CMSWire, 2025*).

The Aflac case is not unusual in being technically sophisticated. By 2026, technical sophistication is widely available. What makes the case strategic is the design discipline behind it — the deliberate placement of AI inside a workflow where its strengths and its limits both translate into business outcomes. That is the question Module 1 is built around.

The Model Is Rarely the Moat

Perhaps the most important shift to recognize as 2026 begins is that the model is rarely the moat. The moat is what the organization builds around the model — the proprietary context the system can reach, the workflow it is embedded in, the loops that turn use into learning, and the trust that lets people rely on the output. Most readers approach AI strategy expecting the central question to be one of model selection: which model, which provider, how to access it. The data and the cases increasingly say otherwise. AI capability is converging across leading providers and becoming

substitutable in important respects, while the organizational capability to deploy AI well remains uneven and slow to build. That asymmetry is where strategy lives.

The executive question has shifted in the three years since generative AI entered most boardroom conversations. In 2023 and 2024, the question was largely one of curiosity: what can these systems do, and where might they apply? In 2025, it became more disciplined: what value are we getting from this investment? In 2026, it has sharpened again — where can durable advantage exist when AI capability is broadly available?

The data behind the shift is real. Stanford's 2026 AI Index puts organizational AI adoption at 88% in surveyed organizations, with generative AI used in at least one business function at 70%.¹ McKinsey's 2025 global survey reports a similar pattern — nearly nine in ten organizations report regular AI use in at least one business function.² Access is widespread; differentiation is not. McKinsey finds that only about a third of organizations are scaling AI across the enterprise, and only 39% report any enterprise-level EBIT impact, with most of those reporting less than 5% of EBIT attributable to AI.² The strategic landscape is no longer "AI haves versus AI have-nots"; it is "AI users versus AI value creators."

This module is about that distinction. Executives do not need to become model engineers. They do need to recognize where AI changes the economics of work, the basis of competition, the speed of organizational learning, the boundaries of products and services, and the defensibility of a firm's advantage.

Reflection prompt as you read. Where is AI already in use in your organization? Where is there clear value? Where is there activity, enthusiasm, or experimentation but no clear enterprise impact yet? Hold these in mind through the rest of the module. The "Bring This to Class" prompt at the end will ask for a more specific version of these answers.

A Business Taxonomy: Analytical, Generative, and Agentic AI

Analytical AI predicts, classifies, detects, forecasts, and optimizes. It is the long-running workhorse of enterprise AI. It answers questions such as: which customers are likely to churn, which transactions look fraudulent, what demand should we expect next quarter, and which maintenance events are likely to happen soon. Analytical AI usually works from structured or semi-structured data and tends to sit closest to measurable operational value.

Generative AI creates, synthesizes, translates, drafts, codes, summarizes, and reasons over unstructured information. Its leverage is high because much executive and professional work is language work — writing, analysis, interpretation, knowledge retrieval, communication, and coordination. Generative AI expands the range of work that software can participate in, but it does not automatically redesign the work.

Agentic AI plans and acts across multiple steps toward a goal. An agent can use tools, retrieve information, call systems, update records, generate outputs, request approval, and continue a workflow. Agentic AI moves the executive question from "What can the model say?" to "What authority should the system have?" A customer-service agent, procurement assistant, research

agent, software engineering agent, or shopping assistant is not a chatbot. It is a partial redesign of a workflow and a partial redistribution of decision rights.

These categories are not mutually exclusive. A modern AI product may use analytical forecasting, generative summarization, and agentic tool use in the same workflow. The strategic point is to ask what kind of work is being changed: prediction and classification, creation and synthesis, or action and workflow execution. That distinction helps avoid two common mistakes. The first is treating every AI opportunity as a generative AI opportunity. The second is assuming agentic AI is a more impressive assistant rather than a new operating-model question.

For strategy purposes, the important issue is not the model family. It is the managerial lever.

AI mode	Primary managerial lever	Typical executive decision
Analytical AI	Better prediction and optimization	Where can better forecasts, classifications, or risk signals change performance?
Generative AI	Faster synthesis and content/code creation	Where does knowledge work slow down because people must read, draft, summarize, or translate too much?
Agentic AI	Delegated action across workflows	What authority can safely move from human-only execution to supervised digital execution?

Consider a customer-service example. Analytical AI can predict which customers are likely to churn or which tickets are likely to escalate. Generative AI can summarize the customer history, draft a response, or translate the answer into the customer's preferred language. Agentic AI can retrieve account information, check policy rules, issue a replacement, update a CRM record, and schedule a follow-up. The strategic question is not which form of AI is more impressive. The strategic question is which combination changes cost, speed, experience, trust, or retention.

A product-innovation example follows the same logic. Analytical AI can identify changes in category demand, pricing sensitivity, or consumer segments. Generative AI can synthesize research, produce concept variations, and draft testing stimuli. Agentic AI can coordinate a concept-testing workflow, route outputs to reviewers, compare results against thresholds, and trigger the next stage. The question is not "Are we using GenAI?" It is "Where in the innovation system does AI shorten learning cycles or improve the quality of choices?"

We treat AI as a strategic, organizational, and leadership problem throughout the course. Technology matters, but value depends on where the technology is placed in the work system.

The 2026 Capability Landscape: Powerful, Jagged, and Substitutable

The pace of AI capability improvement remains rapid by historical standards. Stanford's 2026 AI Index reports continued gains across reasoning, coding, multimodal tasks, and agentic benchmarks. On OSWorld, a benchmark for agents completing real computer tasks, frontier performance has climbed from the low double digits in 2024 to the low-to-mid 80s by mid-2026 — Anthropic reports roughly 83% for Claude Opus 4.8 — though part of the reported gain reflects changes in how the evaluation is run rather than pure capability improvement.³ That is a strategic signal: AI systems are moving from answering questions to completing tasks.

Capability also remains uneven. The phenomenon Andrej Karpathy, formerly of OpenAI and Tesla, has called "jagged intelligence" — models that perform impressively on some advanced tasks while failing on tasks that look simple to humans — appears across the leading 2026 model evaluations.⁴ The empirical pattern matches the "jagged technological frontier" first documented by Dell'Acqua, Mollick, and colleagues in their 2023 BCG/HBS field experiment, in which AI assistance improved consultant performance on tasks inside the frontier and degraded it on tasks just outside.⁵ This unevenness matters for executives because the strategic frontier is not what AI can do in a demo. It is what AI can do reliably, repeatedly, safely, and economically inside the messy constraints of an organization.

The other major 2026 development is model convergence. Stanford reports that top-model performance is increasingly clustered among several leading providers.³ The gap between the best models is often much smaller than the gap between a successful and an unsuccessful deployment. This does not mean all models are interchangeable. They differ in cost, latency, modality, context length, reliability, tooling, privacy, security, and domain performance. It does mean that "we use the best model" is a weak strategy by itself.

For most firms, the durable advantage is unlikely to come from owning a frontier foundation model. It is more likely to come from what sits around the model: the proprietary context the model can access; the workflows it is embedded in; the quality of evaluation and monitoring; the speed with which the organization learns from use; the trust users and customers place in the system; and the ability to redesign work instead of bolting AI onto old processes.

This is the first major 2026 lesson: AI capability is becoming broadly accessible, while AI value remains highly uneven.

Strategic Implications of Model Convergence

Model convergence does not make model choice irrelevant. It means model choice should be treated as one part of a broader enterprise design. A model that performs best on a public benchmark may not be the best model for a regulated workflow, a cost-sensitive high-volume process, a multilingual customer interaction, a private-data environment, or a task that requires low latency. Model selection should not become a proxy for strategy.

For most organizations, the practical question is increasingly a sourcing and architecture question: rely on one model provider, route tasks across several models, use open models for selected workloads, fine-tune or adapt a model for a narrow domain, or buy an application where the model is hidden inside the workflow. Module 3 develops those choices in depth. Module 1 needs only the strategic implication: when models become more substitutable, advantage moves upward into workflow ownership, proprietary context, orchestration, user experience, trust, and learning loops.

This shift also changes the vendor conversation. A demo may show an impressive model output, but the strategic questions sit one layer down. The Vendor Diagnostic in the Field Reference at the end of the module turns those questions into a checklist; for now, the point is that they are strategy questions disguised as technology questions.

The Enterprise Value Gap

The central management challenge is no longer whether executives have heard of AI or whether employees are experimenting with it. The challenge is moving from isolated usage to measurable enterprise value.

McKinsey's 2025 survey captures the pattern. AI use is widespread, but most organizations have not yet embedded it deeply enough into workflows and processes to capture material enterprise-level benefits. The high performers are different: they are more likely to redesign workflows, aim beyond cost reduction, use AI for growth and innovation, track value, and show strong senior-leadership ownership.²

Kimberly-Clark — the consumer products company behind Huggies, Kleenex, and Scott — has framed itself publicly as a "data-driven and AI-First" company. Its Maestro supply-chain orchestration platform has been reported to deliver more than \$20 million in value, and the company's Bengaluru global digital technology center, opened in 2018, has expanded roughly eight-fold over five years (*Kimberly-Clark / PR Newswire, December 2024*). The pattern matches McKinsey's high-performer description: workflows redesigned around AI, named senior ownership, tracked value, and a growth-oriented framing rather than a headcount-reduction one.

This distinction matters in practice. A company can have thousands of AI users and still have little strategic impact. It can issue AI tools to every knowledge worker and still fail to change cycle time, customer experience, product innovation, risk management, or operating margin. Adoption is an input. Value is an outcome.

AI activity	AI value creation
Employees use generic tools	Teams redesign work around AI-enabled processes
Pilots show local promise	Scaled use cases affect operating metrics
The organization celebrates usage counts	Leaders track KPI movement and business outcomes
AI is managed as a technology project	AI is managed as strategy, product, process, risk, and change
Tools are available but optional	AI is embedded in recurring workflows
Benefits are anecdotal	Benefits are linked to cost, revenue, speed, quality, or risk

The implication is not that experimentation is wrong. Early exploration is necessary. The problem is letting experimentation become a substitute for strategy.

MIT Sloan Management Review makes a similar point in its work on "small t" transformations.⁶ Generative AI has not produced the wholesale redesign of major business functions that many expected, but companies are creating real value by moving up a risk slope: starting with tasks, then role-specific workflows, then team processes, and eventually products and customer-facing experiences. That is a useful executive discipline. Do not promise transformation everywhere. Build toward it through teachable, measurable, scalable changes.

Why AI Pilots Stall

Most stalled AI efforts do not fail because the demo was fake. They fail because the demo was not attached to the management system required for value capture. Four patterns recur.

First, pilots target visible tasks rather than important value pools. A tool that drafts emails may be useful, but the financial impact may be small or diffuse. A claims-triage process, pricing workflow, sales-enablement system, or product-testing cycle may be more valuable because it connects to a measurable operating outcome.

Second, pilots leave the workflow intact. Employees are asked to use AI in addition to the old process rather than redesigning the process around new capabilities. This can create more work, not less. If the AI output still has to be copied into another system, checked manually without clear standards, reconciled with old templates, and approved through the same bottlenecks, the organization has added a layer rather than improved a process.

Third, pilots lack benefit ownership. IT, analytics, or innovation teams may build the capability, but the business unit owns the P&L, the customer outcome, and the operating metric. If no executive owns the value case, adoption may become voluntary and measurement may become anecdotal.

Fourth, pilots underinvest in trust. Employees will not rely on AI systems if they cannot understand when the output is grounded, when it needs review, what data it used, or who is accountable for errors. Trust is not a soft issue. In many domains, trust is the adoption mechanism.

These four failure modes give you a practical way to analyze your own organization. The question is not only "Do we have pilots?" It is "Do we have a path from pilot to workflow, from workflow to KPI, and from KPI to strategic value?"

The AI Value Portfolio

To make AI strategy actionable, executives need a portfolio view. Module 1 introduces the course's recurring AI value portfolio. Later modules deepen it with use-case discovery, KPI value chains, ROI logic, governance, and the 90-day action plan.

Efficiency: Doing Existing Work Faster or Cheaper

Efficiency initiatives reduce cost, cycle time, rework, or manual effort. They are usually easiest to measure and often make good early candidates for enterprise adoption. Examples include document processing, call summarization, software testing, invoice matching, claims triage, knowledge retrieval, and internal service-desk automation.

The executive question is: where are we spending scarce human time on repeatable work that can be standardized, assisted, or automated without damaging quality, trust, or accountability?

Truist offers a banking-side example. The bank's AI-enabled integrated receivables platform, launched in February 2026 for commercial and corporate clients, uses machine learning to auto-match payments to invoices — a workflow that has historically required substantial manual reconciliation inside accounts-receivable teams (*Truist press release, PR Newswire, February 3, 2026*). The strategic point is not the model. It is that the work being changed is a measurable operating chokepoint inside a customer relationship that already exists.

Efficiency work should not be dismissed as cost cutting. Well-chosen efficiency projects can build AI literacy, create reusable data and integration foundations, surface governance needs, and establish credibility with finance. The risk is stopping there. If AI strategy becomes only a headcount or

productivity story, the organization may miss opportunities to improve decisions, customer experience, and future growth.

Effectiveness: Doing Work Better

Effectiveness initiatives improve decision quality, customer experience, forecasting accuracy, risk detection, product quality, or employee capability. They are often more strategically valuable than simple automation because they improve the quality of judgment and action.

Examples include sales-enablement systems that synthesize customer context, AI-assisted product research, risk signals for underwriters, demand sensing, pricing recommendations, and executive decision support.

The executive question is: where would better information, faster synthesis, or improved decision framing materially change performance?

Effectiveness projects often require more sophisticated measurement. The benefit may not show up as hours saved. It may show up as better conversion rates, lower loss ratios, fewer stockouts, higher forecast accuracy, faster cycle time, better compliance, or improved customer satisfaction. This is why the course returns later to KPI value chains. Executives must be able to connect AI-enabled work changes to operational and financial indicators.

Innovation: Doing New Things

Innovation initiatives use AI to create new offerings, revenue streams, business models, or customer experiences. These are often higher-risk and harder to measure early, but they may matter most for long-term positioning.

Examples include AI-native professional services, personalized customer journeys, predictive maintenance services, AI-enabled product design, autonomous support experiences, and data products built from proprietary insight.

The executive question is: what customer need, product category, or business model becomes possible because AI can now understand, generate, reason, or act at scale?

Optomed offers a healthcare counterpart. The Finnish company's Aurora AEYE — the first FDA-cleared handheld AI fundus camera, cleared in May 2024 — combines a portable imaging platform with an embedded algorithm to detect more-than-mild diabetic retinopathy in primary care, with reported diagnostic sensitivity of 92–93% and specificity of 89–94% (*Optomed / PR Newswire, May 1, 2024*). The product is strategic because the AI is what makes the workflow possible at the point of primary care. Without the algorithm, the device is incomplete; with it, a specialist diagnostic extends into a generalist setting where it could not previously go.

Innovation projects require a different governance rhythm. They should not be judged only by immediate cost savings, but they also cannot live forever as vague strategic bets. Good innovation governance asks what must be learned by the next milestone. Is there customer demand? Is there a credible path to differentiation? Can the organization deliver the experience reliably? Does the offering improve with use? Can it be priced, sold, supported, and governed?

Portfolio Discipline

The healthiest AI portfolio is rarely all efficiency or all innovation. A portfolio that focuses only on efficiency may produce useful savings while leaving the firm vulnerable to AI-native competitors. A portfolio that focuses only on bold innovation may produce expensive pilots without near-term credibility. Executives need a mix: near-term efficiency plays that build confidence and operating discipline; effectiveness plays that improve important decisions and customer outcomes; and select innovation bets that test where AI could change the basis of competition.

The portfolio is not the business case. It is the starting map. Module 5 develops the measurement, ROI, and stop / pivot / scale logic in more detail.

Durable Advantage When Models Commoditize

Return to the contrarian thesis. The MIT Sloan Management Review provocation "Why AI Will Not Provide Sustainable Competitive Advantage" is useful here.⁷ If AI itself becomes ubiquitous, AI itself cannot be the source of sustained advantage. That does not make AI strategically unimportant. It means the advantage has to be specified more carefully.

Where does it sit?

Generic Model Access Is Rarely a Moat

Using a leading AI model through a vendor API may be useful, but it is rarely defensible. Competitors can often access the same or similar capabilities. If the use case depends only on generic model intelligence, the advantage may be temporary.

Proprietary Context Can Be a Moat

AI becomes more powerful when it has access to the right context: customer histories, product data, research archives, transaction records, contracts, claims, service logs, design history, operating procedures, and domain-specific knowledge. Data alone is not a moat, however. Data becomes strategic when it is unique, current, trusted, legally usable, connected to decisions, and embedded in a workflow.

Colgate-Palmolive offers a useful incumbent example. MIT Sloan Management Review reports that the company uses generative AI and retrieval-augmented generation to synthesize proprietary consumer-insight reports along with trusted external sources.¹⁰ The strategic point is not "Colgate uses a chatbot." It is that AI gives more employees faster access to accumulated market knowledge and helps them participate in innovation cycles.

Georgia-Pacific offers a parallel example in an industrial setting. The company — Koch Industries' pulp, paper, and building products subsidiary — runs a generative AI system called "ChatGP," built on AWS, that surfaces operator and maintenance knowledge from documents, digital files, and what the company calls tribal knowledge across more than 140 facilities (*Georgia-Pacific / AWS case study, 2024*). Roshan Shah, vice president of applied AI and products, has described the deployment as talent-led: the value is not the chatbot but the years of operator know-how the chatbot makes searchable, reusable, and applicable to the next problem (*Roshan Shah, VP of Applied AI and Products, Georgia-Pacific, Manufacturing Digital, 2024*). Like the Colgate example, this is

proprietary context made productive — not by collecting more data but by connecting trusted data to the decisions and routines that already drive the business.

The example also shows why proprietary context is broader than data volume. A company may have many documents but little usable context if the information is fragmented, stale, poorly governed, or disconnected from actual decisions. A smaller but high-quality body of domain knowledge can be more valuable when it is trusted, searchable, permissioned, and connected to a workflow that matters.

Workflow Integration Can Be a Moat

AI that sits outside the flow of work is easy to ignore. AI embedded in a recurring workflow can change how work is performed. The deeper the integration into customer journeys, employee workflows, approvals, systems of record, and feedback loops, the more likely it is to create switching costs and operating advantage.

Amazon's Rufus illustrates this logic in a consumer setting. Amazon describes Rufus as built on model routing, retrieval-augmented generation against the product catalog, customer reviews, community Q&A, customer shopping activity, price history, and agentic features such as adding items to carts or auto-buying at a target price.¹¹ The strategic question is not whether Rufus has the best general model. The question is how Amazon combines proprietary context, customer relationship, workflow integration, and purchasing authority inside a shopping journey. (A fuller diagnostic of Rufus appears in the Field Reference Use-Case Gallery.)

Workflow integration is where many AI strategies become real or fail quietly. If an AI tool is outside the workflow, employees must remember to use it. If it is inside the workflow, the system can shape the next action. A sales assistant that lives outside the CRM may be a helpful writing tool. A sales assistant that reads account history, surfaces risks, drafts next-best actions, updates the CRM, and learns from deal outcomes is closer to an operating capability.

Organizational Learning Can Be a Moat

AI advantage can compound when usage creates feedback that improves the system and the organization. This requires more than storing data. It requires structured learning loops: what did the model recommend, what did the human decide, what happened afterward, was the decision better, and how should the process, prompt, data, policy, or model change next time.

Organizations that build these loops may learn faster than rivals. Organizations that merely deploy tools may not.

AI can accelerate individual work, but the strategic prize is organizational learning. If every AI interaction disappears into an employee's private workflow, the firm learns little. If the organization captures which recommendations were accepted, which were rejected, where humans corrected the system, and which decisions produced better outcomes, AI becomes part of a learning infrastructure rather than a productivity feature.

Trust, Governance, and Reliability Can Be a Moat

In regulated, high-risk, or trust-sensitive settings, reliability is itself strategic. A legal AI product, medical workflow, financial-risk system, or procurement agent does not win because it is creative. It wins because users trust its evidence, controls, permissions, auditability, and human oversight.

Trust can become a source of differentiation when it lets an organization move faster with evidence, reduce adoption friction, satisfy customers or regulators, and create confidence that competitors cannot easily copy. The Aflac example at the start of the module sits here as much as anywhere — the design discipline that kept human judgment in the higher-stakes claims is what made the automation in the lower-stakes claims credible to customers and regulators.

Human Judgment Becomes More Important, Not Less

If AI becomes broadly available, human judgment does not become less important. It may become more important. When more firms have access to similar tools, differentiation shifts toward better questions, sharper strategic choices, creative product thinking, faster organizational learning, better implementation, and disciplined refusal to automate the wrong work.

This is a useful place to connect back to the MIT SMR provocation. If AI becomes a general-purpose capability that lifts many firms, then uniquely human and organizational capabilities become more visible sources of advantage. The question is not whether humans or AI are better. The question is how the organization combines machine scale with human taste, judgment, empathy, domain knowledge, and strategic courage.

A Six-Question Test

The Field Reference contains the Durable AI Advantage Test as a standalone tool. The short version, for use as you read on:

1. Is the context proprietary, trusted, and legally usable?
2. Is the AI embedded in a recurring workflow or customer journey?
3. Does use create a feedback loop that improves performance?
4. Does the system change decision speed, quality, cost, risk, or customer experience?
5. Would switching away be operationally painful for users or customers?
6. Can the value be measured through a KPI chain?

If the answer to most of these is no, the initiative may still be useful, but it is probably not a durable strategic moat.

AI-Native Competition and Incumbent Reinvention

AI-native companies matter because they do not simply add AI to an existing process. They often redesign the product, service, or workflow around AI from the start.

Cursor (developed by Anysphere, which SpaceX announced it would acquire in June 2026 at a reported valuation near \$60 billion) is one example. Its strategic significance is not that it uses language models to write code. The deeper point is that it changes the developer workflow. It sits inside the coding environment, participates in editing and debugging, and shapes how developers

think through software creation.¹² The moat question is whether the product owns the workflow deeply enough that switching becomes painful.

Harvey is a parallel example in legal work. Its strategic significance is not document summarization. It is the attempt to embed AI into high-value legal workflows where domain trust, source grounding, client confidentiality, and professional accountability matter.¹³ The moat question is whether it can combine domain-specific context, customer trust, workflow integration, and quality control better than generic AI tools.

Perplexity is a third example, in search and knowledge work. Its strategic significance is not answering questions. It is challenging the interface, economics, and user expectations of search by combining retrieval, synthesis, citations, and conversational exploration.¹⁴ The moat question is whether it can build a trusted user habit and information workflow that incumbent search providers cannot easily absorb.

Lockheed Martin offers an instructive incumbent counterpoint. The company's internal AI Factory — its enterprise MLOps and generative AI platform — supports tens of thousands of internal users and thousands of internal AI agents built on its Genesis platform. In December 2024, Lockheed Martin spun out Astris AI as a subsidiary to commercialize that internal capability for the broader defense industrial base and for high-assurance commercial customers (*Lockheed Martin press release, December 16, 2024*). Astris AI's customers and partners now include Dell Technologies, HPE, Oracle Cloud, NVIDIA, Meta, and several smaller defense suppliers. The Lockheed Martin AI Center team behind the platform was led at launch by Mike Baylor, Lockheed Martin's vice president and chief digital and AI officer, and Greg Forrest, now Lockheed Martin's vice president of AI foundations and commercialization (*Lockheed Martin press release, December 16, 2024; Astris AI board listing, 2026*).

The strategic logic is worth pausing on. Lockheed Martin did not respond to AI-native competition by becoming a startup. It identified an internal capability — high-assurance MLOps developed under defense-grade oversight — that other organizations needed and could not easily build, and turned it into an external offering. The capability was an incumbent asset before it was an AI product. The AI commercialization made the asset visible.

The danger for incumbents is using AI only to make today's process marginally cheaper while AI-native competitors redesign the process entirely. The opportunity is the inverse, and it rests on a forty-year-old finding. David Teece's research on who profits from innovation showed that when the innovation itself is easy to imitate, returns flow to whoever holds the complementary assets the innovation needs to reach customers: distribution, brand, installed relationships, regulatory position, specialized capability.¹⁵ Jay Barney's factor-market logic sharpens the point: a resource any competitor can buy at market price cannot differentiate.¹⁶ Model access is now exactly that resource. Lovable, the Stockholm platform behind much of the "vibe coding" wave, crossed \$400 million in annual recurring revenue in February 2026 with 146 employees (*TechCrunch, March 2026*). That is a demonstration of how cheap building has become, and of how thin a moat building alone provides. The assets your organization already has — proprietary data and institutional knowledge, installed customer relationships, distribution and trust, operational scale, regulatory experience, domain expertise, embedded customer or employee workflows — are the complementary assets of the AI era. The strategic question is which of them become more valuable

when AI is applied to them, or applied around them. The strongest incumbents will not simply defend the old model. They will use AI to make their existing assets work in new ways. The strongest challengers will not simply wrap a model. They will own a painful workflow, improve with use, and build trust in a domain where generic tools are not enough.

Drafting Your Strategic Thesis

Before class, draft one sentence using this template:

Our organization should use AI to [change this source of value] by [embedding AI in this workflow / customer journey / decision process], creating advantage through [proprietary context, workflow integration, trust, learning speed, or business-model innovation].

Three calibration examples:

- Our organization should use AI to improve commercial decision speed by embedding customer and product intelligence into account-planning workflows, creating advantage through proprietary customer context and faster organizational learning.
- Our organization should use AI to reduce claims cycle time and improve customer trust by embedding document intelligence into claims triage, creating advantage through proprietary claims history, operational workflow integration, and measurable service quality.
- Our organization should use AI to accelerate new product development by connecting consumer insight archives, market signals, and concept-testing workflows, creating advantage through proprietary research context and faster innovation cycles.

The thesis is intentionally rough. It will be refined across the course as you work through executive AI literacy, use-case discovery, stack and sourcing, governance, economics, capability acceleration, and implementation planning.

Bring This to Class

Come prepared with three things:

1. **Three AI initiatives from your organization, placed into the Value Portfolio.** Where do they fall — efficiency, effectiveness, innovation, or hybrid? What is missing? Does the portfolio overweight productivity, contain ambitious experiments without near-term value evidence, or include enough disciplined innovation to test how AI could change the basis of competition?
2. **Your draft Strategic Thesis sentence,** built from the template above.
3. **A two-minute answer to a board-level question:** does your organization need its own AI assistant — a Rufus-equivalent for your domain? Be ready to argue either side. The case for or against will hinge less on whether the technology is impressive than on what proprietary context, workflow integration, customer relationship, and learning loop your organization could bring to it.

A Final Word Before Class

Module 1 has worked through a taxonomy, a portfolio, six potential moats, and a diagnostic. These are useful tactics, and you will use them through the rest of the course. They should not, however, obscure the higher-altitude point. AI strategy in 2026 is not a question about which model to license or which vendor to standardize on. It is a question about where in your organization the technology should sit — which workflows it joins, which judgment it supports, which decisions it accelerates, which trust it earns. The tactics are how that question gets answered. The rest of the course exists to develop the answer. The first move is yours.

Part 2: Field Reference

The Field Reference is the leave-behind. Use it during the course to scaffold board memos, vendor conversations, and pilot reviews; keep it after the course as an executive reference. Each component is self-contained — flip directly to a table, prompt list, or worksheet without reading surrounding prose.

Decision Tables

AI Mode → Managerial Lever

AI mode	Primary managerial lever	Typical executive decision
Analytical AI	Better prediction and optimization	Where can better forecasts, classifications, or risk signals change performance?
Generative AI	Faster synthesis and content/ code creation	Where does knowledge work slow because people must read, draft, summarize, or translate too much?
Agentic AI	Delegated action across workflows	What authority can safely move from human-only execution to supervised digital execution?

AI Activity vs. AI Value Creation

AI activity	AI value creation
Employees use generic tools	Teams redesign work around AI-enabled processes
Pilots show local promise	Scaled use cases affect operating metrics
The organization celebrates usage counts	Leaders track KPI movement and business outcomes
AI is managed as a technology project	AI is managed as strategy, product, process, risk, and change
Tools are available but optional	AI is embedded in recurring workflows
Benefits are anecdotal	Benefits are linked to cost, revenue, speed, quality, or risk

AI-Native Challenger vs. Incumbent

Question	AI-native challenger	Incumbent
Starting point	Designs around AI-enabled workflow from day one	Begins with existing products, processes, customers, and constraints
Advantage	Speed, focus, clean workflow design, fewer legacy constraints	Data, trust, distribution, relationships, domain knowledge, operating scale
Risk	Thin moat if built only on generic model capability	Incrementalism and slow internal adoption
Strategic test	Can it own a workflow deeply enough to become hard to replace?	Can it redesign the workflow before a challenger does?

Durable AI Advantage Test

For any AI initiative, score 1–5 against each question. Total range 6–30; treat 22+ as a candidate for "scale," 14–21 as "redesign," below 14 as "pause" or "defer."

Diagnostic question	Score (1–5)
Uses proprietary context competitors cannot easily buy	
Embedded in a recurring workflow or customer journey	
Generates feedback that improves the system	
Changes decision speed, quality, cost, risk, or experience	
Increases switching costs or deepens trust	
Can be linked to a KPI chain	

Vendor Diagnostic

Use this checklist when evaluating any AI vendor — model provider, application vendor, or systems-integration partner. Demos answer the wrong questions. These are the right ones.

1. **Context.** What context does the system use? How is the context curated, refreshed, and governed? What is in it, and what is explicitly excluded?
2. **Evaluation.** How is accuracy evaluated? What benchmarks, what production sampling, what error categories, what threshold for action? Who reviews the results?
3. **Exception handling.** How does the system handle questions or cases outside its intended scope? Does it escalate, decline, hedge, or hallucinate?
4. **Learning.** How does the system learn from human corrections? Are corrections captured, analyzed, and folded back into the system, or discarded?
5. **Integration.** How does the system integrate with our systems of record? What writes back, what does not, and through what controls?
6. **Substitution.** What happens if we switch models or providers? Is the application portable, or are we locked in by the underlying model?
7. **Value measurement.** What value metric improves after deployment? Who owns that metric, and what is the baseline before the deployment starts?

A vendor that cannot answer these without redirecting to a feature demo is selling you a technology, not an enterprise capability.

Do / Don't Quick Reference

Do	Don't
Target pilots at value pools where the operating metric is named and owned	Target pilots only at visible tasks because they are easy to demo
Redesign the workflow around the AI capability	Bolt the AI onto the existing workflow and call it transformation

Do	Don't
Assign business-unit ownership of the value case before building	Run the pilot from IT or innovation alone with no accountable P&L owner
Invest in trust mechanics — grounding, citations, review paths, audit	Treat trust as a soft issue to be addressed after rollout
Build learning loops that capture corrections and outcomes	Let every AI interaction disappear into private user workflows
Treat model selection as one decision inside an enterprise design	Treat model selection as a substitute for strategy
Pair near-term efficiency wins with disciplined innovation bets	Run a portfolio that is all efficiency, all innovation, or all theater
Quote dated, sourced data; name your evidence	Quote "many companies struggle" without a denominator

Cautionary Tales

Air Canada and the Liable Chatbot (February 2024)

A grieving Air Canada customer asked the company's chatbot whether bereavement fares could be applied retroactively. The chatbot said yes; airline policy said no. When the customer was denied the discount, he sued. Air Canada argued that the chatbot was "a separate legal entity responsible for its own actions." A Canadian tribunal disagreed and ordered the airline to honor the chatbot's commitment. **The lesson:** statements your AI agents make to customers, partners, or regulators may bind your organization. Workflow integration and authority go together; granting authority without grounding the system in your actual policies is a liability you cannot disclaim.

McDonald's, IBM, and the Drive-Thru Pullback (June 2024)

After three years of partnership across more than 100 drive-thru locations, McDonald's ended its IBM-developed AI voice-ordering pilot. Public footage of order errors — McNugget orders that climbed past 200 items, mistaken bacon on ice cream — circulated widely on social media. The technology problem was not the underlying model. It was a deployment that did not adapt the workflow (background noise, accent diversity, menu complexity) to the system's actual reliability profile. **The lesson:** "the demo worked" is not a proof of production readiness. A pilot that targets a visible task without redesigning the workflow around the system's strengths and limits will fail visibly.

Workbook Prompt: Apply the Durable AI Advantage Test

Run this as a 30-minute exercise with three to five colleagues from across your organization (operations, technology, finance, line-of-business leadership).

1. **Pick one AI initiative** — proposed or in flight — at your organization. Write down a one-sentence description.
2. **Score the initiative** against the six Durable AI Advantage Test questions (above). One scorer per question; rotate. Total the score.

3. **Discuss the lowest two scores.** What would have to change about the initiative to raise each by two points? What would that cost?
4. **Decide a recommendation:** scale, redesign, pause, or defer.
5. **Write a three-sentence memo** to a hypothetical board: what the initiative is, what the score is, what the recommendation is, and what would change the recommendation.

Repeat with two more initiatives. The pattern across three initiatives often surfaces more about your organization's AI strategy than any single deep dive.

Use-Case Gallery: Amazon Rufus

Amazon describes Rufus as a shopping assistant powered by generative and agentic AI. It uses a mix of models accessed through Amazon Bedrock, including Anthropic Claude, Amazon Nova, and a custom Amazon model. It draws on Amazon's product catalog, customer reviews, community Q&A, customer shopping activity, product knowledge graphs, retrieval-augmented generation against selected sources, price history, and customer context. As of Amazon's Q4 FY26 earnings (February 2026), Rufus had reached approximately 300 million customers, and customers who used it while shopping were more likely to make a purchase during that shopping trip.¹¹

The board question. Is Rufus primarily an efficiency tool, an effectiveness tool, an innovation tool, or a strategic moat?

A good answer surfaces that it may be all four:

- *Efficiency:* reduces customer effort and support burden.
- *Effectiveness:* improves product discovery and decision quality.
- *Innovation:* creates a more agentic shopping experience.
- *Moat:* deepens Amazon's customer context, workflow integration, and purchasing relationship.

The counterargument. If other retailers can license similar models, add conversational shopping, and connect their own catalogs, Rufus is not automatically defensible. Defensibility depends on Amazon's proprietary context, scale, customer habit, data feedback loops, broader commerce and fulfillment ecosystem, and ability to act on behalf of the customer.

The translation question for your organization. A board member asks whether your company — large retailer, bank, insurer, manufacturer, or professional-services firm — needs its own Rufus-equivalent. A strong answer avoids imitation for its own sake. It identifies the workflow or customer journey where an AI assistant could create measurable value, the proprietary context the firm can bring to the assistant, the risks of granting it authority, and the KPI that would prove progress. The point is not to copy Amazon. The point is to translate the moat logic into your own context.

Part 3: Card Specification

The AI Strategic Thesis Card

The AI Strategic Thesis card is drafted in session; there is no post-class submission. The five elements below define a complete card. Teams use them in the Moat Diagnostic, and students on the Blog 3 Transform path use them afterward to develop the card via the Module 8 Capstone Artifact Kit.

Required elements:

1. **Strategic opportunity:** What business problem or opportunity does AI address?
2. **Value pillar:** Is it efficiency, effectiveness, innovation, or hybrid?
3. **Possible durable moat:** What could make this hard to copy?
4. **Primary uncertainty:** What could prevent value capture?
5. **One metric:** What KPI would prove progress?

Capstone build (Transform path): the card becomes the opening strategic thesis of the final executive AI strategy memo. Expect to revise it as Modules 3, 4, 5, and 7 add stack, governance, economics, and capability evidence.

Next-module handoff: Module 2 takes this strategic thesis into daily executive practice by asking how leaders delegate AI-assisted work, define context, and review output responsibly.

Works Cited and Sources

Inline-cited sources (executives and named-company examples):

- Shelia Anderson, then CIO, Aflac (departed for Unum Group, May 2025). MIT Sloan Management Review *Me, Myself, and AI* podcast, "Continuous Learning With AI: Aflac's Shelia Anderson," 2023.
- Keith Farley, Vice President of Innovation Operations, Aflac. CMSWire, "Human First: How Aflac Combines AI With Authentic Connections," 2025.
- Roshan Shah, Vice President of Applied AI and Products, Georgia-Pacific. Manufacturing Digital, "KONE & Georgia-Pacific, elevators & paper, gen AI & results," 2024; AWS Case Study, "Georgia-Pacific Optimizes Operator Efficiency Using Generative AI on AWS," 2024.
- Lockheed Martin press release, "Lockheed Martin Launches Astris AI to Bring Defense-Grade AI Capabilities to Industry," December 16, 2024.
- Truist press release, "Truist launches AI-enabled receivables platform to bring simplicity, speed, safety and smarter experiences to commercial and corporate clients," PR Newswire, February 3, 2026.
- Kimberly-Clark, "India's Thriving Tech Hub Fuels Kimberly-Clark's Digital Transformation," PR Newswire, December 26, 2024.

- Optomed, "Optomed Releases First FDA-Cleared Handheld AI Fundus Camera for Detection of More Than Mild Diabetic Retinopathy," PR Newswire, May 1, 2024.
 - Air Canada chatbot tribunal ruling, *Moffatt v. Air Canada*, British Columbia Civil Resolution Tribunal, February 2024.
 - McDonald's / IBM AI drive-thru pilot termination, multiple outlets reporting June 2024.
1. Stanford HAI. *The 2026 AI Index Report* (overview, economy, technical performance). <https://hai.stanford.edu/ai-index/2026-ai-index-report>
 2. McKinsey & Company. "The State of AI in 2025: Agents, Innovation, and Transformation." <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai/>
 3. Stanford HAI. "Technical Performance — The 2026 AI Index Report." <https://hai.stanford.edu/ai-index/2026-ai-index-report/technical-performance>; Anthropic, "Introducing Claude Opus 4.8" and Claude Opus 4.8 System Card, May 28, 2026 (OSWorld-Verified ~83%, with the prior Opus 4.7 score restated under a revised evaluation harness). <https://www.anthropic.com/news/claude-opus-4-8>
 4. Andrej Karpathy, "Jagged intelligence" post, X (formerly Twitter), July 25, 2024.
 5. Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2023). "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality." Harvard Business School Working Paper No. 24-013, September 15, 2023.
 6. MIT Sloan Management Review. "Generate Value From GenAI With 'Small t' Transformations." <https://sloanreview.mit.edu/article/generate-value-from-gen-ai-with-small-t-transformations/>
 7. MIT Sloan Management Review. "Why AI Will Not Provide Sustainable Competitive Advantage." <https://sloanreview.mit.edu/article/why-ai-will-not-provide-sustainable-competitive-advantage/>
 8. BCG. "Leading in the Age of AI Agents: Managing the Machines That Manage Themselves." <https://www.bcg.com/publications/2025/machines-that-manage-themselves>
 9. MIT Sloan Management Review. "How to Scale GenAI in the Workplace." <https://sloanreview.mit.edu/article/how-to-scale-genai-in-the-workplace/>
 10. MIT Sloan Management Review. "The GenAI Focus Shifts to Innovation at Colgate-Palmolive." <https://sloanreview.mit.edu/article/the-genai-focus-shifts-to-innovation-at-colgate-palmolive/>
 11. Amazon. "Amazon's Next-Gen AI Assistant for Shopping Is Now Even Smarter, More Capable, and More Helpful." <https://www.aboutamazon.com/news/retail/amazon-rufus-ai-assistant-personalized-shopping-features>
 12. Anysphere, the company behind Cursor, was valued at \$29.3B in November 2025 (Series D, \$2.3B raised). A reported ~\$50B round in April 2026 did not close; on June 16, 2026, SpaceX announced an all-stock acquisition of Anysphere at a reported ~\$60B valuation, placing Cursor under its xAI subsidiary. Annualized revenue reached \$2B in February 2026, with press reports placing it above \$3B by mid-2026. Deal terms are press-reported and the transaction had not closed as of July 1, 2026. TechCrunch, "Sources: Cursor in talks to raise \$2B+ at \$50B valuation as enterprise growth surges," April 17, 2026, <https://techcrunch.com/2026/04/17/sources-cursor-in-talks-to-raise-2b-at-50b-valuation-as-enterprise-growth-surges/>; press coverage of the SpaceX announcement, June 16, 2026.
 13. Harvey raised \$200M at an \$11B valuation in March 2026, co-led by GIC and Sequoia, bringing total funding to more than \$1B and reflecting deployment across more than 25,000 custom legal agents. Harvey, "Harvey Raises at \$11 Billion Valuation to Scale Agents Across Law Firms and Enterprises," March 25, 2026, <https://www.harvey.ai/blog/harvey-raises-at-dollar11-billion-valuation-to-scale-agents-across-law-firms-and-enterprises>; Bloomberg and CNBC coverage, March 25, 2026.
 14. Harvard Business School. "Perplexity: Redefining Search." <https://www.hbs.edu/faculty/Pages/item.aspx?num=67198>
 15. David J. Teece, "Profiting from Technological Innovation: Implications for Integration, Collaboration, Licensing and Public Policy," *Research Policy* 15, no. 6 (1986): 285-305.
 16. Jay B. Barney, "Strategic Factor Markets: Expectations, Luck, and Business Strategy," *Management Science* 32, no. 10 (1986): 1231-1241.

Copyright, reuse, and citation

© 2026 Gerald C. Kane. This module is licensed under the Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0), <https://creativecommons.org/licenses/by-nc/4.0/>.

You may share and adapt it for noncommercial purposes, provided you credit Gerald C. Kane as the author, link to the license, and indicate any changes you made. You may not present this material, or adaptations of it, as your own work, and you may not sell it or use it in paid training or consulting without written permission.

Cite as: Kane, G. C. (2026). *Module 1: AI Strategy Landscape*. EMBA-7450, Information Technology and Strategy, Executive MBA, Terry College of Business, University of Georgia. <https://www.profkane.com/course>

Permissions and commercial use: jerry@profkane.com