
EMBA-7450 · INFORMATION TECHNOLOGY AND STRATEGY · BONUS MODULE

AI Strategy Capstone Artifact Kit

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: Capstone artifact kit. The module is not a teaching module. It carries the discipline argument for the optional executive AI strategy memo, the assembly map for the seven prior course artifacts, and the field-reference templates that translate the course frameworks into a senior-team decision document.

Module Learning Objective

By the end of this module, you will be able to assemble the seven course artifacts you have built across Modules 1-7 into a 5-7 page executive AI strategy memo for your own organization or a chosen company, and you will be able to evaluate a memo of this kind against the criteria a senior team would apply. For Fall 2026, the assembled memo is offered as an optional capstone deliverable. Students who do not undertake the optional capstone will still find the Field Reference templates useful for executive AI strategy work after the course.

Module Role in the Course

This module sits after Module 6 (the integration capstone with the in-class board simulation) and Module 7 (capability acceleration). Earlier modules each produced a named artifact: the AI Strategic Thesis memo (Module 1), the AI Delegation Card (Module 2), the Stack and Sourcing Appendix (Module 3), the AI Governance Card (Module 4), the AI Business Case (Module 5), the Board Strategic Response (Module 6), and the AI Capability Acceleration Plan (Module 7). The Bonus Module's job is to give those seven artifacts a coherent home as one decision document.

The Pre-Class Read carries the discipline argument and the assembly map. The Field Reference carries the templates. Part 3 carries the evaluation rubric.

Part 1: Pre-Class Read

BOTTOM LINE. A board-ready AI strategy memo is a decision document, not a literature review or a vendor pitch. The test of whether it works is whether a senior team could use it to make a decision without the author in the room. The discipline lives in specificity — one named basis of competition, one priority workflow, one accountable owner, one falsifiable success signal, one credible 90-day path — rather than in breadth of frameworks reviewed. Across the course, students have already built the raw material. The work of the capstone memo is not to produce new content. It is to assemble the existing artifacts into one document that earns the senior team's time.

Opening Vignette: What Discipline Sounds Like at the Senior Level

When Derek Waldron, Chief Analytics Officer of JPMorgan Chase, described the firm's AI program at the 2025 Investor Day and in subsequent McKinsey coverage, he did not talk about transformation, leverage, or unlocking value. He named numbers and cadence. Roughly 200,000 employees on LLM Suite — approximately two-thirds of the workforce. About \$1.5 billion in fraud and operations savings in calendar year 2024. Investment banker pitch-deck generation cut from hours to roughly thirty seconds. A 10-20% developer productivity uplift on coding assistants. An eight-week platform release cycle that matches the rate at which models, vendors, and use cases change underneath them (*Derek Waldron, Chief Analytics Officer, JPMorgan Chase, McKinsey interview, October 2025; JPMorgan 2025 Investor Day disclosures*).¹

Waldron's framing is what discipline sounds like at the senior level. There is one named platform (LLM Suite), one quantified workforce reach (200,000), one specific dollar figure tied to a named year (\$1.5B in CY2024), one named productivity claim with a bounded range (10-20% on coding), one operating cadence (eight weeks). A board could use that framing to make decisions: approve the next round of capability investment against the disclosed savings, evaluate the next agentic-coding investment against the developer-productivity claim, hold the eight-week cadence as the cycle against which annual planning gets stress-tested. A literature review of generative-AI use cases in financial services, however well written, would not give a board the same purchase.

The work of your final memo is to produce the equivalent — for your own organization or a chosen one — at a level of specificity that earns the same kind of decision conversation.

The Memo Discipline Most Strategy Documents Skip

Perhaps the most important discipline in writing a final AI strategy memo is also the one most often skipped: choosing what *not* to include. The temptation in a 5-7 page document is to demonstrate breadth — frameworks reviewed, technologies surveyed, competitors compared, regulations summarized, governance structures mapped. A memo that does this earns a high reading grade and a low decision grade. The senior reader closes the document with a sense of the topic and no sense of what to do next.

The disciplined memo trades breadth for argument. It names one basis of competition that has actually shifted, one priority workflow where the shift can be tested, one authority level that the AI system has earned, one set of risks that controls actually mitigate, one KPI value chain that connects technical activity to financial and strategic outcomes, one operating model that names the accountable owner, and one 90-day path that the organization could plausibly start tomorrow. Breadth is sacrificed deliberately. The memo's argument is that this is the move worth making first; everything else is acknowledged briefly and held for later.

The seven course artifacts already encode this discipline. The capstone memo's job is to assemble them.

The Course Artifact Assembly Map

You should not treat the final memo as a fresh assignment. It is the assembly and refinement of the artifacts you have built across the course. The map below points back to where each artifact was produced; revisit your own work as the raw material before drafting.

Source artifact	Where you built it	What it contributes to the memo
AI Strategic Thesis memo	Module 1 in-class card	Frames where AI changes the basis of competition for the organization and which value pillar matters. Becomes the memo's strategic thesis section. Revise after Modules 3, 4, 5, and 7 added stack, governance, economics, and capability evidence.
AI Delegation Card	Module 2 in-class card	Defines the workflow, context boundary, context quality check, human review gate, harness needs, and authority level for one recurring executive workflow. Becomes the spine of the prioritized use case section.
Stack and Sourcing Appendix	Module 3 in-class card	Explains technical commitments, context engineering choices, harness design, sourcing choices, observability, and portability. Becomes the sourcing and stack implications section.
AI Governance Card	Module 4 in-class card	Names risk tier, controls, assurance evidence, oversight, auditability, and incident readiness. Becomes the risk and governance section.
AI Business Case	Module 5 in-class card	Connects the use case to the KPI value chain, TCO/FinOps, harness economics, cost per successful task, benefits ownership, and stop/pivot/scale criteria. Becomes the KPI value chain and economic argument.
Board Strategic Response	Module 6 in-class team exercise	Rehearsed executive synthesis under competitive pressure in the Module 6 board scenario. Becomes the model for how the memo's argument should hold together under board questioning.
AI Capability Acceleration Plan	Module 7 in-class card	Shows how the strategy becomes executable through managers, champions, cascade learning, context and harness ownership, governance feedback, and first-90-day evidence. Becomes the operating model and 90-day plan sections.

The final memo should revise, compress, and integrate these artifacts. It should not paste them together.

What Good Looks Like

The Memo Component standards below name the section structure of a memo that earns a senior team's time. The Field Reference templates support each section; the standards are what the rubric in Part 3 actually evaluates against.

Memo Component	Target Length	What Good Looks Like
Executive summary	0.5 page	Names the recommendation, the strategic rationale, and the top decisions required. A senior reader who reads only this page knows what to approve.
Strategic AI thesis	0.5 page	Explains how AI changes the basis of competition for the organization. Names a specific competitive shift, not a general transformation.
AI value portfolio	1 page	5–8 initiatives balanced across efficiency, effectiveness, and innovation. The portfolio shows the option set the recommendation came from, not just the chosen path.
Prioritized use case	1 page	Specific workflow, named users, data sources, context quality concerns, agent authority level, harness needs, value logic, and launch path. The level of specificity is workflow, not function.
Sourcing and stack implications	0.5–1 page	Build / buy / partner decision tied to speed, control, cost, data, context ownership, harness requirements, and risk. The recommendation is one sentence; the rationale is the page.
Risk and governance	1 page	Top risks, mitigations, evidence, owners, and escalation path. Risks are named at the level of mechanism (goal drift, data leakage, regulatory exposure), not at the level of category.
KPI value chain and 90-day plan	1–1.5 pages	Metrics from technical to operational to financial to strategic, with cost per successful task and eval/monitoring cadence. The 90-day plan names actions, owners, and decision gates, not aspirations.

A Final Word Before You Start

The kit that follows looks like a stack of templates. It is. The discipline, however, is not in the templates — it is in the writing. Each template is a forcing function: a set of fields you have to answer concretely before you can move on. The templates work because they expose the places where breadth would otherwise replace specificity. A blank "primary users / stakeholders" field cannot be filled with "the organization"; a blank "agent authority level" cannot be filled with "appropriate"; a blank "evidence of progress" cannot be filled with "ongoing." The templates are scaffolds; the memo is yours.

Use them.

Part 2: Field Reference

The templates below translate the seven course artifacts into the section structure of the final memo. They are organized in roughly the order a senior team would read them: cover sheet first, thesis next, portfolio, prioritized use case, authority, sourcing, context and harness, risk, KPI chain, operating model, 90 days, board briefing.

Template 1: Board Memo Cover Sheet

Field	Response
Organization / business unit	
Decision requested from leadership	
Strategic AI thesis in one sentence	
Primary value objective	Efficiency / Effectiveness / Innovation / Balanced portfolio
Recommended first use case	
Agent authority level	Assistant / Copilot / Workflow Agent / Supervised Digital Worker / Multi-Agent Process
Sourcing recommendation	Build / Buy / Partner / Hybrid
Top risk requiring executive attention	
90-day success signal	
Context engineering issue requiring leadership attention	
Harness or operating control requiring leadership attention	
Cost per successful task metric	

Template 2: Strategic AI Thesis

Complete the prompts below before filling out the portfolio. The thesis should be sharp enough to guide tradeoffs.

Prompt	Response
Where is AI most likely to change the basis of competition in this industry?	
Which assets or capabilities are most defensible for this organization?	
Where could an AI-native competitor attack our economics or customer relationship?	
Where could customer, supplier, or employee agents change behavior?	
What is our strategic posture for the next 12 months?	Defend / Catch up / Differentiate / Reinvent

Template 3: AI Value Portfolio

List 5–8 candidate initiatives. A strong portfolio includes near-term efficiency, capability-building effectiveness, and at least one disciplined innovation option.

Initiative	Value Type	Workflow / Owner	Expected Business KPI	Risk / Dependency	Decision
	Efficiency / Effectiveness / Innovation				Fund / Explore / Hold / Stop
	Efficiency / Effectiveness / Innovation				Fund / Explore / Hold / Stop
	Efficiency / Effectiveness / Innovation				Fund / Explore / Hold / Stop
	Efficiency / Effectiveness / Innovation				Fund / Explore / Hold / Stop
	Efficiency / Effectiveness / Innovation				Fund / Explore / Hold / Stop
	Efficiency / Effectiveness / Innovation				Fund / Explore / Hold / Stop
	Efficiency / Effectiveness / Innovation				Fund / Explore / Hold / Stop

Template 4: Prioritized Use Case Charter

The Use Case Charter is the strategic frame for the prioritized initiative. A CEO or business unit leader could complete it without a CIO in the room. Template 7 covers the technical and operational complement.

Charter Element	Response
Business problem	
Target workflow	
Current pain point or constraint	
Primary users / stakeholders	
Agent role and authority level	
Human oversight and escalation path	
Minimum viable launch scope	
Scale path after proof of value	
Stop / pivot / scale criteria	
Named owner for delivery	

Template 5: Agent Authority and Controls

Authority Level	Use in This Strategy?	Decision Rights	Controls Required	Owner
Assistant		Drafts, summarizes, or analyzes only	Human review; source checking; privacy boundaries; context-quality check	
Copilot		Recommends actions inside a human-led workflow	Approval gate; confidence signals; logged recommendations; eval checklist	
Workflow Agent		Completes bounded tasks in approved systems	Least privilege; tool limits; escalation rules; test sets; observable tool calls	
Supervised Digital Worker		Acts across systems within defined thresholds	Audit trail; runtime monitoring; kill switch; exception handling; cost limits	
Multi-Agent Process		Coordinates multiple agents across an end-to-end process segment	Risk tiering; red team; operating model owner; incident response; harness assurance	

Template 6: Build / Buy / Partner Decision

Criterion	Build	Buy	Partner	Implication for Recommendation
Speed to value				
Control over IP and data				
Integration complexity				
Talent availability				
Operating cost and AI FinOps				
Vendor lock-in / exit path				
Regulatory and assurance needs				
Strategic differentiation				
Context ownership and freshness				
Harness control and observability				
Eval, monitoring, and cost transparency				

Recommendation sentence: We recommend [build / buy / partner / hybrid] because the strategy requires [speed / control / cost / risk reduction / differentiation], and the binding constraint is [data, talent, integration, regulation, capital, or time].

Template 7: Context, Data, and Harness Implications

The Context, Data, and Harness template is the technical and operational complement to Template 4. It surfaces the questions a CIO, data leader, or risk leader would ask before approving the use case for production.

Question	Response
What proprietary data or context makes this use case defensible?	
Where is the data today, and who owns it?	
What quality, access, privacy, or lineage gaps must be fixed?	
Does the use case need RAG, structured data integration, tools, external systems, memory, or cached context?	
What evaluation, monitoring, and observability are required?	
What cloud, model, or infrastructure dependency is strategically material?	
What are the context engineering choices: write, select, compress, isolate, and cache?	
Which knowledge files, source repositories, or project instructions must stay fresh?	
Who owns context quality, knowledge-file freshness, and source retirement?	
What harness components are required: system prompt, tools, retrieval, memory, output structure, evals, monitoring, and orchestration?	
Who owns harness assurance: prompt and project libraries, eval routines, tool settings, retirement decisions?	
What is the expected cost per successful task, and what could cause it to rise?	

Template 8: Use-Case Risk Register

Risk	Likelihood / Impact	Owner	Mitigation	Evidence / Monitoring	Escalation Trigger
Goal drift / unsafe autonomy ²					
Tool misuse or excessive permissions ²					
Data leakage / privacy ³					
Bias / unfair impact ³					
Regulatory noncompliance ⁴					
Vendor lock-in / infrastructure dependency					
Workforce trust / adoption failure					
Context failure: stale, missing, contradictory, poisoned, over-permissioned, or excessive context ³					
Harness failure: weak evals, missing monitoring, tool overreach, no rollback, or unclear ownership ³					
Knowledge-file or project-library decay					

Template 9: KPI Value Chain

The purpose of the KPI value chain is to prevent teams from celebrating technical activity that does not move business results.

Layer	Metric	Baseline	Target	Owner	Review Cadence
Technical	Accuracy / groundedness / task completion / latency / context length / cache-hit rate / cost per successful task				
Operational	Cycle time / first-contact resolution / exception rate / throughput / quality				
Financial	Cost-to-serve / revenue lift / margin / working capital / avoided loss				
Strategic	Retention / differentiation / share gain / customer trust / option value				
Assurance	Eval pass rate / monitoring alerts / review reversal rate / incident count				

Template 10: AI Operating Model

Operating Model Decision	Recommendation
Who owns enterprise AI strategy?	
What belongs in a central AI office or platform team?	
What is embedded in business units?	
How do product, data, technology, legal, HR, finance, and risk coordinate?	
How does a use case move from idea to production?	
What steering committee reviews portfolio, risk, value, and capacity?	
What is the decision cadence for fund / stop / pivot / scale?	
What evidence must be reviewed before increasing agent authority?	
Who owns context quality, knowledge-file freshness, and source retirement?	
Who owns prompt / project libraries and assistant portfolio hygiene?	
Who runs eval routines and decides when quality drift requires changes?	
Who monitors cost per successful task and cache economics?	

Template 11: First 90 Days

Phase	Objective	Actions	Evidence of Progress
Days 1-30: Assess and align	Create a shared baseline and decision frame.	Inventory priority use cases; align executive thesis; classify risks; identify one workflow for deeper design.	Leadership agreement on thesis, portfolio, owner, and risk tiering.
Days 31-60: Architect and initiate	Design the launch path and controls.	Select sourcing path; define data needs; build KPI chain; create test/eval plan; assign operating model roles.	Approved use case charter, governance path, and launch plan.
Days 61-90: Launch and learn	Begin controlled deployment and value tracking.	Run limited pilot; monitor technical and business KPIs; collect user feedback; update risk register; decide fund / stop / pivot / scale.	First measurable workflow result and decision on next investment.

Template 12: Five-Minute Board Briefing

The briefing is the spoken complement to the written memo. It answers six questions in roughly thirty seconds each.

1. *The strategic threat or opportunity.* What changed and why it matters now.
2. *The recommended strategic move.* Where to focus first and why.
3. *The value logic.* How the initiative moves operational and financial KPIs.
4. *The operating model.* Who owns it, how decisions will be governed, and how context and harness quality will be maintained.
5. *The risk posture.* What could go wrong, including context or harness failures, and how management will know early.
6. *The ask.* What the board or executive team should approve today.

Source Guidance

Sources support the memo where they strengthen executive judgment. The memo is not a literature review. Use citations selectively, attached to claims that need them.

The Works Cited list below carries seven external sources plus one corpus citation. Each one supports a particular kind of claim:

- **Capability trends and adoption.** When the memo references the state of AI capability, deployment rates, or industry-level adoption patterns, the Stanford HAI 2026 AI Index Report⁵ is the cleanest source. McKinsey's 2025 *State of AI*⁶ and *The Big Rethink*⁷ cover enterprise adoption and operating-model patterns specifically.
- **Infrastructure constraints.** When the memo argues that compute, energy, or data center capacity is a strategic constraint for the recommended path, the IEA's *Energy and AI*⁸ is the source.
- **Regulatory timing.** When the memo references AI Act, AI literacy, GPAI, or high-risk AI rule timing, the European Commission AI Act implementation timeline⁴ is the source. For U.S. policy framing, the White House *America's AI Action Plan*⁹ is the source.
- **Risk and assurance language.** When the memo names risk categories, controls, or assurance evidence, the NIST AI RMF Generative AI Profile³ supports lifecycle and governance language; the OWASP Top 10 for Agentic Applications² supports agent-specific risk language (goal hijacking, identity abuse, tool misuse, prompt injection, rogue autonomous behavior).
- **Named-company precedent.** When the memo cites a specific company's AI program as analog or contrast, the citation should be a real published source — earnings call, executive interview, peer-reviewed case, or credible secondary reporting. The opening vignette uses Waldron at JPMorgan Chase¹ as one example.

The rubric in Part 3 treats citation discipline as part of strategic clarity. A memo that makes a regulatory-timing claim without citing the timeline, an infrastructure-constraint claim without

citing the data, or a named-company precedent without sourcing it, is read as breadth-without-discipline.

Part 3: Evaluation Rubric

The rubric evaluates a memo against eight criteria. Each criterion names what a strong submission does and the most common failure mode that produces a weaker submission.

Criterion	Strong Submission	Common Failure Mode
Strategic clarity	Names the real basis of competition for the organization and the specific way AI shifts it.	Claims the basis of competition has changed without specifying how, or treats AI itself as the competitive shift rather than the underlying advantage AI enables.
Portfolio discipline	Balances near-term efficiency, capability-building effectiveness, and at least one disciplined innovation option. The portfolio shows the option set the recommendation came from.	Lists initiatives at the same value type (typically all efficiency), or presents a single recommended initiative without showing what was considered and rejected.
Use-case specificity	Defines the prioritized use case at workflow level, with named users, data sources, decision rights, and a credible launch path.	Defines the use case at function level (e.g., "customer service") rather than workflow level (e.g., "tier-1 billing inquiry resolution at first contact"); leaves users, data, or decision rights unnamed.
Governance and risk	Names risks at the level of mechanism, with credible controls, owners, and escalation triggers.	Names risks at the level of category (e.g., "regulatory risk") without naming the mechanism, the control that addresses it, or the trigger that escalates.
KPI chain	Connects technical or agent activity to operational, financial, and strategic outcomes, with cost per successful task and review cadence.	Reports technical metrics (accuracy, latency, completion rate) without connecting them to operational or financial outcomes; treats technical performance as the value claim.
Operating model	Clarifies ownership, cadence, cross-functional coordination, and decision rights for both the strategy and the context/harness work.	Names a steering committee without naming the accountable owner; leaves context quality and harness assurance ownership unstated; describes coordination without describing decision rights.
Executive usefulness	A senior team could plausibly use the memo to make a decision without the author in the room.	Describes options without recommending one, or qualifies the recommendation so heavily that no decision is on the table; the memo is a survey rather than a decision document.
Context and harness design	Names the context sources, freshness ownership, harness components, eval routines, and cost per successful task.	Treats context and harness as implementation detail; names the AI capability without naming what makes it dependable in operation.

Works Cited and Sources

1. Derek Waldron, Chief Analytics Officer, JPMorgan Chase, McKinsey interview, October 2025; JPMorgan Chase 2025 Investor Day disclosures; *American Banker*, "Innovation of the Year: JPMorgan's LLM Suite," June 2025.
2. OWASP GenAI Security Project. *Top 10 for Agentic Applications for 2026*. <https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/>
3. NIST. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence>
4. European Commission AI Act Service Desk. *AI Act Implementation Timeline*. <https://ai-act-service-desk.ec.europa.eu/en/ai-act/timeline/timeline-implementation-eu-ai-act>

5. Stanford HAI. *2026 AI Index Report*. <https://hai.stanford.edu/ai-index/2026-ai-index-report>
6. McKinsey. *The State of AI in 2025: Agents, Innovation, and Transformation*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai/>
7. McKinsey. *The Big Rethink: An Agenda for Thriving in the Agentic Age*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-big-rethink-an-agenda-for-thriving-in-the-agentic-age>
8. International Energy Agency. *Energy and AI*. <https://www.iea.org/reports/energy-and-ai/executive-summary>
9. The White House. *America's AI Action Plan*. <https://www.whitehouse.gov/releases/2025/07/white-house-unveils-americas-ai-action-plan/>

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). *Bonus Module: AI Strategy Capstone Artifact Kit*. 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