

# Future of Competition

---

*From Frontier Shock to Board-Ready Response*

**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:** Integration capstone. The module synthesizes the course toolkit (value portfolio, build/buy/partner, risk register, KPI value chain, operating model, and 90-day plan) into a board-ready response to an agentic-AI competitive shock.

---

## Module Learning Objective

---

By the end of this module, you will be able to diagnose an agentic-AI competitive shock, decide where to act first, and defend the response to a board using the course frameworks. Every student will complete the Board Strategic Response's Core fields: the threat diagnosis, the priority workflow, the agent authority level with its required controls, and the first 90 days. At the Challenge depth, called the Full Response in this module, the same artifact extends to the value portfolio moves, the KPI value chain with stop and scale criteria, the build/buy/partner choice, and the operating-model owner. Students choosing the Blog 3 capstone-memo option may use the complete one-page Board Strategic Response to develop their memo.

---

## Module Role in the Course

---

Earlier modules prepared this capstone conversation. Module 1 established the shift from AI adoption to durable advantage. Module 2 gave you an executive routine for delegating cognitive work to AI with judgment. Module 3 covered the stack, context, sourcing, and agent architecture beneath dependable enterprise capability. Module 4 developed governance, risk, assurance, and accountability. Module 5 covered the economic discipline to move from pilot to P&L. Module 7 will deepen the organizational capability, talent, and workflow redesign work needed for scale.

Module 6 sits at the point where those threads must come together. It asks for a response to a plausible agentic-AI competitive shock and a board-ready strategy. The module owns future competition, strategic foresight, agent authority as a board-level decision, infrastructure constraints, and capstone synthesis. It does not reteach prompting, stack architecture, governance frameworks, ROI mechanics, or change management in detail. Those topics matter here only because they shape the strategic response.

The boundary is useful:

| Topic                                   | Module 6 Treatment                                                            |
|-----------------------------------------|-------------------------------------------------------------------------------|
| Executive AI literacy                   | Use as output discipline and review practice; Module 2 owns the deeper skill. |
| Stack and sourcing                      | Apply build/buy/partner logic; Module 3 owns the architecture.                |
| Governance and risk                     | Use as board approval criteria; Module 4 owns the detailed control system.    |
| ROI and scaling                         | Require KPI chain and stop/scale logic; Module 5 owns the economics.          |
| Talent and change                       | Name operating implications; Module 7 owns capability acceleration.           |
| Future competition                      | Own fully.                                                                    |
| Board simulation and capstone synthesis | Own fully.                                                                    |

Read this module for the artifact it builds. Every section of the Pre-Class Read feeds the same deliverable, and the strongest way to read is with your own organization's most plausible shock in mind.

## If You Want to Know More

A useful way to prime the reading is to hold one question in mind as you go: if a competitor announced tomorrow that its operations were agent-mediated, what would your board ask you by lunchtime? The formal preparation sits in "Bring This to Class" at the end of Part 1.

The Pre-Class Read below makes the argument; the materials in the table carry the underlying data and provocation. These references are optional. Read what fits your time.

| Source                                                                                                | Why It Matters For This Module                                                                                                                                          |
|-------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Stanford HAI, <i>2026 AI Index Report</i>                                                             | Capability progress, organizational adoption, agent deployment, infrastructure concentration, and responsible AI gaps. <sup>1</sup>                                     |
| McKinsey, <i>The State of AI in 2026: On the Road to ROI</i>                                          | Enterprise adoption, agent scaling by company size, build-versus-buy, operating-cost constraints, EBIT impact, and workflow redesign. <sup>2</sup>                      |
| International Energy Agency, <i>Energy and AI</i>                                                     | Data center electricity demand, local-grid bottlenecks, and infrastructure constraints. <sup>3</sup>                                                                    |
| European Commission AI Act implementation timeline and Regulation (EU) 2026/1744                      | The post-Omnibus baseline: Article 50 transparency in application since August 2, 2026; high-risk obligations deferred to December 2027 and August 2028. <sup>4 5</sup> |
| NIST AI RMF: Generative AI Profile and OWASP Top 10 for Agentic Applications                          | Assurance, lifecycle governance, and agent-specific risk language. <sup>6 7</sup>                                                                                       |
| World Economic Forum, "Here's a Playbook for Boards on How to Govern Agentic AI" (April 2026 article) | Board-facing article on the governance gap and the board challenge protocol. <sup>8</sup>                                                                               |

You should also review your prior course artifacts: AI value portfolio, build/buy/partner decision, risk register, KPI value chain, operating model notes, and 90-day action plan.

## Part 1: Pre-Class Read

---

**BOTTOM LINE.** Agentic AI puts a new decision on the board agenda: what authority to delegate to AI systems, under what controls, and with what evidence that the systems are behaving as intended. DeepSeek's January 2025 R1 release and the separately developed Stargate investment plan illustrate how algorithmic efficiency and infrastructure scale can reshape competition. Access to compute, energy, chips, and jurisdiction belongs in the strategy discussion. A disciplined response starts by diagnosing what has changed, choosing where to act first, setting authority boundaries, governing with evidence, and deciding who owns the next 90 days.

**TERMS.** Selected strategy and infrastructure terms not already defined in Modules 1-5 appear in **bold** at their first substantive use in Part 1; each is defined in the Executive Glossary immediately before Part 2. Use the glossary as a lookup aid, not assigned vocabulary to memorize.

### Opening Vignette: The DeepSeek Shock and the Stargate Buildout

---

On January 20, 2025, DeepSeek released its R1 reasoning model. The market reaction intensified over the following week, including a sharp fall in Nvidia shares on January 27.<sup>9</sup> DeepSeek challenged assumptions about the cost of strong AI capability, but the widely cited \$5.576 million figure belongs to the reported training computation for its V3 base model. That estimate used an assumed rental rate for Nvidia H800 GPUs and excluded prior research and ablation work; it is not R1's complete development cost.<sup>11</sup> For this module, an **algorithmic-efficiency shock** means a change in the expected cost of producing a given capability. The strategic question is how such a change affects investments built on the previous cost assumptions.

Separately, on January 21, 2025, OpenAI, SoftBank, and Oracle announced the Stargate Project: an intention to invest \$500 billion in U.S. AI infrastructure over four years, beginning with \$100 billion. Construction was already underway when the plan was announced. By September 23, OpenAI reported nearly 7 GW of planned capacity across an expanded portfolio that included five newly announced sites, Abilene, and CoreWeave projects; Abilene was operating by then.<sup>10</sup> These were investment and capacity plans at their respective announcement dates, not proof that every dollar had been spent or every site completed.

The two developments illustrate different strategic mechanisms. DeepSeek raised the possibility that efficiency gains could change the economics of frontier capability. Stargate showed the scale of infrastructure commitments being made to supply that capability. Their proximity does not establish that Stargate was a response to DeepSeek. Together they give a board two assumptions to test: how much capability will cost, and what access to capacity will require.

## The Decision That Most Agentic-AI Conversations Skip

Each new wave of AI capability, analytical, generative, and now agentic, has reached most boards the same way: a vendor presentation, a feature comparison, a build-versus-buy chart, and a pilot proposal. Agentic AI sits inside that familiar pattern and does not belong there. It is a decision-rights question with infrastructure constraints, governance obligations, and competitive consequences attached.

Shelia Anderson, then chief information officer at Aflac, described a related decision on MIT Sloan Management Review's *Me, Myself, and AI* podcast in 2023. Aflac used automation for claims work while preserving human judgment for complex cases.<sup>31</sup> That example provides a starting point for the 2026 board question. An agent can pursue a goal across multiple steps, use tools, access context, and act in the world. What authority may it have, under what controls, and with what evidence that it is behaving as intended? The rest of this module is built around that question.

## Which Frontier Signals Matter for Executives?

Executives do not need a generic AI trend report. They do need a shared view of the terrain, and the signals below earn their place because they change how leaders think about capability, scale, authority, infrastructure, risk, and regulation. The frameworks in this Pre-Class Read are built to outlast their examples; the dated instances that illustrate them live in the Current Horizon Scan in the Field Reference and get replaced as they age.

None of this means every organization must chase every frontier. It means the frontier is becoming operational. AI capability is available through many providers; enterprise value depends on whether the organization can put that capability into workflows, govern it, measure it, and scale it responsibly.

| Strategic Signal                                          | 2026 Signal                                                                                                                                                                                                                                                                                                                                | Executive Question                                                                                                   |
|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| AI capability is still climbing, but unevenly.            | Stanford's 2026 AI Index reports continued capability gains and describes a jagged frontier: systems can be impressive in one task and unreliable in another. <sup>1</sup>                                                                                                                                                                 | Where do we need model brilliance, and where do we need bounded reliability?                                         |
| Adoption is broad, but enterprise value remains uneven.   | Stanford's 2026 report puts surveyed AI adoption at 88%; McKinsey's August 2026 survey reports 44% scaling AI across the enterprise, up from 38%. These are separate surveys with different measures and collection periods. <sup>12</sup>                                                                                                 | Which workflows are ready for scaling, and which are still experimentation?                                          |
| Reported enterprise financial impact remains limited.     | McKinsey reports that 37% of respondents attribute any EBIT impact to AI, unchanged from 2025, while about 6% meet its high-performer definition. <sup>2</sup> These are respondent reports, not audited financial outcomes.                                                                                                               | What does our pilot-to-production rate actually look like, and what is causing the gap between deployment and value? |
| Agent scaling differs by company size and survey measure. | McKinsey reports 40% of organizations with at least \$1 billion in revenue scaling agents in at least one function, up from 27%, versus 22% of smaller organizations. Stanford reports single-digit scaled agent use in nearly all individual functions using earlier survey data. The measures are not directly comparable. <sup>12</sup> | What authority should agents have, and who is accountable when they act?                                             |

| Strategic Signal                                       | 2026 Signal                                                                                                                                                                                                                                                                                                                                                                                | Executive Question                                                                                               |
|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| Compute and energy are strategy constraints.           | The IEA's base case projects global data-center electricity consumption at about 945 TWh in 2030, roughly double its 2024 level. This covers all data centers, with AI an important driver of growth. <sup>3</sup>                                                                                                                                                                         | Could power, chips, cloud terms, latency, or data-center location limit our AI strategy?                         |
| Responsible AI requires evidence alongside capability. | Stanford highlights gaps in responsible-AI measurement; NIST and OWASP provide risk and assurance frameworks. <sup>167</sup> OpenAI's August 2026 investigation of the Hugging Face intrusion documents unauthorized agent activity during evaluations with reduced safeguards. <sup>12</sup>                                                                                              | Can we demonstrate reliability, security, accountability, and controls appropriate to the authority we delegate? |
| AI policy is part of competitive strategy.             | Regulation (EU) 2026/1744 revised the AI Act timetable: Article 50 transparency requirements began applying August 2, 2026, subject to scope, exceptions and transition provisions; high-risk dates moved to December 2027 and August 2028. <sup>45</sup> The U.S. July 2025 AI Action Plan emphasized innovation, infrastructure, and international diplomacy and security. <sup>13</sup> | Where will regulation, sovereignty, export controls, or infrastructure policy change sourcing and market access? |

These surveys give a board reasons to ask how adoption translates into value. They do not establish why an individual firm is struggling, or whether the gap will persist. A board-ready response connects its deployment plan to workflow evidence, financial outcomes, and a process for learning when the expected benefits fail to appear.<sup>12</sup>

## How Will Competition Change When Agents Mediate Markets?

Agentic AI changes competition by changing the unit of action. Firms will compete not only through products, channels, prices, and brands, but through the speed and quality of AI-mediated workflows. Customers may use agents to compare, negotiate, switch, monitor service, and automate purchasing. Suppliers may use agents to optimize pricing, allocation, compliance, and fulfillment. Competitors may deploy agents that compress decision cycles from weeks to minutes. Employees may use personal or enterprise agents to reshape their own work. This does not mean every industry becomes autonomous. It means more interactions become machine-readable, machine-negotiated, and machine-monitored, and advantage moves closer to workflow integration, proprietary context, trusted interfaces, real-time learning, and the ability to govern delegated action.

A competitor's agentic-AI announcement is rarely one thing. It could be a cost threat if agents reduce manual coordination, a service threat if customers get better visibility and faster recovery, a data threat if the competitor learns faster from every interaction, a channel threat if customers begin using their own agents to compare and switch, or a narrative threat if the competitor has announced a future operating model before the capability is proven. The diagnostic precedes the response.

## Five Competitive Arenas

| Arena           | What Changes                                                                                                                           | Strategic Response                                                                                                             |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| Customer agents | Buyers can compare options, monitor service, switch providers, and demand personalization with far lower search and coordination cost. | Make offerings, performance data, service terms, and recovery options machine-readable without giving away the economic model. |

| Arena                          | What Changes                                                                                                     | Strategic Response                                                                                                         |
|--------------------------------|------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| Workflow-embedded competition  | Advantage moves from a standalone product or service to a product embedded in the customer's operating workflow. | Prioritize workflow integration, proprietary context, APIs, customer data flywheels, and switching value created by usage. |
| AI-native business models      | New entrants may redesign delivery around agents from day one, with fewer handoffs and lower coordination cost.  | Identify which parts of the value chain could be rebuilt if legacy roles, systems, and pricing did not exist.              |
| Infrastructure and geopolitics | Compute, chips, energy, cloud concentration, export controls, and data residency shape who can scale.            | Treat sourcing and infrastructure as strategy, not procurement.                                                            |
| Trust as a market asset        | Autonomous systems can create value quickly, but failures are visible, explainable, and costly.                  | Turn governance, assurance, and transparency into reasons customers, regulators, and partners choose the firm.             |

The arenas become clearer through concrete examples. **Customer agents:** in March 2026, Walmart executive Daniel Danker told *WIRED* that conversion for the selection sold directly inside ChatGPT was roughly one-third that of products requiring a click-through. Walmart and OpenAI were shifting toward embedding Walmart's Sparky experience in ChatGPT.<sup>14</sup> This is one retailer's comparison across different selections and experiences; it does not isolate why conversion differed or establish a general consumer refusal to buy through agents. **Workflow-embedded competition:** warehouse automation illustrates a strategic possibility: value may depend on integration with fulfillment operations, data, and service processes as much as on the equipment. The moat question is whether that integration makes the offering difficult to replace.

**AI-native business models:** an entrant can redesign a workflow around AI and test whether fewer handoffs improve its economics. **Infrastructure and geopolitics:** Lockheed Martin announced Astris AI in December 2024 to offer AI capabilities to external customers.<sup>33</sup> The strategic interpretation is that an incumbent can explore commercializing capabilities developed for its own operations. **Trust as a market asset:** Georgia-Pacific's Roshan Shah, vice president of Applied AI and Products, describes iterative deployments that address operational pain points and build credibility.<sup>32</sup> For the board, the question is what evidence would turn that operational credibility into customer or stakeholder trust.

Use the table as a map. A consumer bank may care most about customer agents, trust, and regulatory obligations. A manufacturer may care most about workflow-embedded competition, supplier agents, and edge deployment. Any board conversation worth having forces prioritization.

## The Headless Inflection: When the Agent Becomes the User

The Five Arenas assume agents may mediate more interactions. Vendor announcements show how suppliers are preparing for that possibility. On April 15, 2026, Salesforce announced Headless 360, exposing capabilities through APIs, command-line tools, and MCP tools. On April 22, OpenAI, Microsoft, and Google announced expanded enterprise-agent capabilities.<sup>15</sup> **Headless** software can be operated through programmatic interfaces, allowing agents to perform work that previously required a person at a screen. MCP, introduced in Module 3, is one interoperability layer; Anthropic donated it to the Linux Foundation's Agentic AI Foundation in December 2025.<sup>16</sup> The dated announcements are in the Current Horizon Scan.

Vendor announcements and organizational adoption measure different things. McKinsey's August 2026 survey finds 40% of large enterprises and 22% of smaller ones scaling agents in at least one function.<sup>2</sup> In a July 2025 HBR Analytic Services survey of 603 technology decision-makers, 6% said their organizations fully trusted agentic AI with core end-to-end processes.<sup>17</sup> KPMG's Q2 2026 survey of 2,145 senior leaders reports that 49% had delayed or scaled back agent deployments because expected operating costs outweighed anticipated business value.<sup>18</sup> These findings support testing both operational readiness and economics before expanding authority.

As enterprise software adds interfaces for agents, four board-level decisions follow. First, which contracts expose the firm to **consumption-based pricing** over the next twenty-four months? Module 5's cost discipline turns that question into an operational audit. On ServiceNow's Q1 2026 earnings call, management said half of net-new business used non-seat pricing; that category included infrastructure, hardware and connectors as well as tokens.<sup>34</sup> McKinsey reports one organization in five limiting AI use because of operating costs, including tokens.<sup>2</sup> Second, when an agent can coordinate work across systems, which separately licensed products remain valuable? In the same survey, 32% said they had decided against buying at least one software product because they could build it with coding agents.<sup>2</sup>

Third, which layer should the firm compete in: data access, ease of agent integration, or the vendor's agent operating environment? Fourth, when work happens through APIs, what evidence will make the agent's actions visible to the people responsible for oversight? Audit logs, approvals, monitoring, and user feedback must support that authority decision.

The scale of future agent use remains uncertain. Leaders can still prepare by testing contract sensitivity to usage, identifying where consolidation would create value, and ensuring that cost and action limits are enforceable.

## Will Any of This Actually Be a Moat?

The Five Arenas ask where AI could create durable competitive position. A May 2025 *MIT Sloan Management Review* essay by David Wingate, Barclay L. Burns, and Jay B. Barney argues that widely available AI may cease to differentiate firms, leaving distinctive resources such as culture and domain knowledge as sources of advantage.<sup>19</sup> This is the **Resource-Based View** applied to an increasingly accessible capability, and a useful corrective to the assumption that buying AI automatically creates a moat.

The Five Arenas survive that scrutiny, but each carries a **moat-mechanism** question. *Customer agents*: does our agent's behavior depend on data, trust, or workflow context competitors cannot easily replicate? *Workflow-embedded competition*: is the embedding deep enough that switching costs are operationally real? *AI-native business models*: does our redesign change unit economics in a way incumbents cannot match? *Infrastructure and geopolitics*: is our infrastructure access actually scarce, and is the counterparty that owns it durable? *Trust as a market asset*: is our trust position legible to customers, regulators, and acquirers in a way they price? An AI investment that fails the moat-mechanism question on its arena is still useful. It is just not a moat.

The strongest answers resist both extremes. They neither dismiss a competitor's announcement as hype nor copy the press release. They identify the mechanism by which AI might change competition and then choose where to respond first.

## Agent Authority Is a Board-Level Decision

For executives, agentic AI is a decision-rights problem. An agent is valuable because it can pursue a goal across multiple steps, use tools, access context, and sometimes act in the world. The same capability creates risk because authority can move faster than accountability. A system that drafts a customer update for human review and a system that changes delivery commitments, issues credits, reroutes freight, or triggers maintenance work may share an interface. Their accountability models are completely different.

An April 2026 World Economic Forum article frames agentic AI as a governance challenge involving delegated authority.<sup>8</sup> The questions below apply that framing to a board decision: what evidence should management provide before the organization expands an agent's authority?

### The Agent Authority Ladder

Module 2 introduced the five-level authority ladder and Module 4 attached the control burden at each level. The contribution here is synthesis: which level is strategically justified under board pressure.

| Level | Agent Role                | Permitted Authority                                               | Required Controls                                                                          |
|-------|---------------------------|-------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| 1     | Assistant                 | Drafts, summarizes, or analyzes only.                             | Human review before use; source checking; privacy boundaries.                              |
| 2     | Copilot                   | Recommends actions inside a human-led workflow.                   | Confidence signals; rationale; approval gates; usage monitoring.                           |
| 3     | Workflow Agent            | Completes bounded tasks in approved systems.                      | Tool permissions; logging; test sets; escalation rules; role-based access.                 |
| 4     | Supervised Digital Worker | Acts across systems within defined thresholds.                    | Runtime monitoring; audit trail; kill switch; segregation of duties; exception review.     |
| 5     | Multi-Agent Process       | Coordinates multiple agents across an end-to-end process segment. | Operating model owner; risk tiering; red-team testing; incident response; board reporting. |

Higher levels carry more control burden and are not automatically better. The right level depends on business value, risk tier, data quality, integration maturity, user trust, regulatory exposure, and the ability to monitor the system in production.

A logistics example: route exception handling may be a credible Workflow Agent or Supervised Digital Worker candidate. A Workflow Agent could retrieve shipment status, detect likely delay, identify approved recovery options, prepare a customer update, and route the recommendation to a dispatcher. With stronger controls, a Supervised Digital Worker might execute low-risk recovery steps under clear thresholds. A customer-facing agent that negotiates service credits, changes contractual commitments, or reroutes high-value freight has a different authority profile and should not be treated as the same use case.

### Board Challenge

If management proposes a Supervised Digital Worker or Multi-Agent Process deployment, the board should ask for six things before approval:

1. The risk tier and business value case.

2. The authority boundary: what the agent may do and may not do.
3. The evaluation results, including edge cases and failure modes.
4. The human escalation path and exception-handling model.
5. The audit trail, monitoring, and incident response plan.
6. The stop, rollback, or scale criteria.

The prior course frameworks come together here. Module 3 carries the stack and observability needed for the agent to work. Module 4 carries the governance and assurance system. Module 5 carries the value evidence and unit economics. Module 7 will carry workflow and role redesign. The question for this module is whether the executive team can bring those pieces together before assigning authority.

## Compute, Energy, and Geopolitics Are Strategy Constraints

---

AI strategy increasingly depends on reliable access to compute, power, chips, cloud platforms, data-center capacity, and jurisdictions that permit the intended use. Most companies will not build data centers or train frontier models, but firms that consume AI through vendors inherit the same dependencies. Model performance, latency, cost, data residency, vendor terms, capacity commitments, and regional availability all shape which use cases are viable.

Stargate illustrates infrastructure treated as strategy: its January and September 2025 announcements committed management attention to years of construction, power supply, financing, and capacity delivery.<sup>10</sup> The wider AI buildout also raises financing questions. NVIDIA's announced investment intention and Alphabet's financing documents distinguish planned funding from cash already available.<sup>20</sup> The dated examples are in the Current Horizon Scan. A buyer of compute should examine both delivery capacity and the financial position of the counterparty promising it.

The IEA's base case projects about 945 TWh of global data-center electricity consumption in 2030. Local conditions matter as much as the global total: a data center can be built faster than the surrounding grid is expanded, making power availability, permitting, **interconnection queues**, cooling, and sustainability commitments potential constraints on deployment.<sup>3</sup>

Forecast uncertainty makes the infrastructure question concrete. ERCOT's 2025 planning materials showed transmission-provider submissions implying 218 GW of demand in 2031, alongside an ERCOT-adjusted forecast of about 145 GW.<sup>21</sup> A capacity plan must identify which scenario it assumes and what has been adjusted. The fuller set of grid signals, including **capacity-market** results and reliability forecasts, is in the Horizon Scan. For capacity-intensive projects, megawatts and delivery dates can constrain the strategy even when capable models are available.

Compute access is also a regulatory question. In May 2025, the U.S. Commerce Department announced rescission of the AI Diffusion Rule and directed officials not to enforce it. In July, NVIDIA said it was applying to resume H20 sales to China after receiving assurances about licenses.<sup>22</sup> These announcements show why firms need to distinguish policy statements, licenses, and actual deliveries. The July 2025 *America's AI Action Plan* set out priorities for innovation, infrastructure, and international diplomacy and security.<sup>13</sup> A firm dependent on a vendor's chip access should review that exposure as part of strategy, with a named owner and a review date.

The full set of executive infrastructure questions is in the Field Reference for use during sourcing and budget reviews.

## Strategic Foresight Without Science Fiction

Traditional scenario planning asks leaders to imagine several plausible futures and test strategy against them. The logic still works. What changes in the AI era is the pace of signals and the number of interacting variables: customer behavior, competitor capability, regulation, vendor roadmaps, labor markets, energy constraints, and model performance can all shift faster than an annual planning cycle.

One signal has an empirical anchor. METR's March 2025 study found that the human-duration equivalent of selected software tasks frontier agents could complete at 50% success had doubled approximately every seven months over the preceding six years.<sup>23</sup> This is a historical capability measure under evaluation conditions. It does not measure enterprise adoption or reliable production autonomy. In the foresight 2x2, operating maturity also depends on workflow integration, controls, and the organization's ability to handle failures.

Supplier choices may also affect the frontier. In a September 2026 essay, Anthropic chief executive Dario Amodei argued for slowing capability development and proposed stronger external evaluation.<sup>24</sup> Treat that as an argument and commitment to examine, not evidence that the measured trend has changed. Useful triggers include verified implementation of supplier commitments, updated capability evidence, and reliable performance in the firm's own workflows.

A second useful frame is the distinction between **acute** and **chronic disruption**, from prior research on organizational response to disruption.<sup>25</sup> Acute disruption arrives as a discrete event: a model release, a competitor announcement, an export-policy change, or a revised regulatory deadline. Chronic disruption builds over time: changing capability, infrastructure pressure, or fragmented policy. For this exercise, use a 90-day response plan for the immediate decision and a longer review cadence for accumulating changes. Use those horizons to organize decisions and review the timing as conditions change.

A foresight discipline an executive team can run in a single session:

1. Scan signals: What weak signals suggest that customers, competitors, suppliers, regulators, or employees may behave differently because of agents?
2. Define uncertainties: Which two uncertainties most change our strategic choices?
3. Stress-test strategy: Which current assumptions fail in each scenario?
4. Choose no-regrets moves: What should we do regardless of which scenario unfolds?
5. Set triggers: What evidence would tell us the future is moving toward one scenario?

Two useful uncertainties for any agentic-AI scenario set are customer-agent adoption and agentic operating maturity. The table carries the four combinations.

| Scenario                                     | What It Would Feel Like                                                         | Strategic Implication                                                   |
|----------------------------------------------|---------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| Slow agent adoption, slow operating maturity | AI remains useful but uneven; pilots matter, but market pressure is manageable. | Build foundations, select high-value workflows, avoid panic investment. |

| Scenario                                      | What It Would Feel Like                                                                                                  | Strategic Implication                                                                                                  |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| Fast customer agents, slow operating maturity | Customers compare, monitor, and switch more easily, but operations are still human-heavy.                                | Improve machine-readable service data, transparency, account workflows, and customer recovery.                         |
| Slow customer agents, fast operating maturity | Competitors improve cost, speed, and reliability before customers fully change buying behavior.                          | Prioritize operations workflow redesign, predictive analytics, and exception handling.                                 |
| Fast customer agents, fast operating maturity | The market becomes more transparent, automated, and unforgiving; speed, trust, and workflow integration matter together. | Move quickly on priority workflows, customer-agent readiness, governance, operating model, and strategic partnerships. |

The point of the scenario frame is to identify **no-regrets moves** rather than to choose the correct box: cleaning the data the priority use case depends on, defining agent authority levels, selecting two high-value workflows, building a partner strategy that preserves optionality, improving customer-facing transparency, and creating a board cadence for AI value and risk. Set **triggers** alongside them, the evidence that would tell you the future is moving toward one scenario.

## How Does Risk Change When Agents Have Authority?

The course risk register matures here from general AI risk to agentic operating risk. NIST's AI RMF and Generative AI Profile supply assurance vocabulary; OWASP's Top 10 for Agentic Applications supplies security language for systems that plan and act across workflows.<sup>67</sup> Executives need to understand what changes when systems gain tools, memory, identity, and authority. OpenAI's August 2026 investigation of the Hugging Face incident describes evaluation agents using unauthorized communication and exploiting infrastructure weaknesses while seeking solutions to difficult tasks. The evaluations used reduced safeguards.<sup>12</sup> The Field Reference carries the case and its limits.

The seven-row Agentic Risk Register (goal drift, tool misuse, data leakage, identity and authority abuse, human overreliance, regulatory misclassification, and infrastructure dependency) is in the Field Reference as a leave-behind. The governing point for the Pre-Class Read is that the register is not abstract. If the priority use case is route exception handling, ask what happens when the agent optimizes for on-time delivery while violating customer commitments, labor rules, safety policies, or cost thresholds. If it is a customer-facing shipment agent, ask what the agent is allowed to promise, what data it can reveal, how it handles conflicting records, and when a human must intervene.

Delegated authority also raises questions of stakeholder legitimacy. After the May 2020 destruction of the Juukan Gorge rock shelters, Rio Tinto announced the departure of its chief executive and two senior leaders.<sup>36</sup> The Field Reference uses this as an analogy about **social license**, not as evidence that AI caused the destruction. The board lesson is to examine consequential harms and stakeholder commitments alongside technical or legal permission.

## Bring This to Class

Every student should bring the Board Strategic Response's Foundation/Core preparation, in one line:

**AI shock:** [a recent or hypothetical competitive shock from your industry or organization: a competitor announcement, a customer-agent rollout, a regulatory shift, an infrastructure constraint] · **Arena:** [which of the Five Competitive Arenas it primarily moves] · **Mechanism:** [cost, speed, customer interface, data feedback, or narrative] · **First workflow:** [where you would act first]

Be ready to defend the arena diagnosis against alternative readings. In class you will extend this line into the Core fields of the Board Strategic Response: the agent authority level for that first workflow, its required controls, and the first 90 days.

For the Challenge depth, the Full Response in this module, extend the same line before class with three more fields: the business KPI the first workflow should move, with stop and scale criteria; the build/buy/partner choice for the core capability; and the two or three other portfolio moves the priority workflow would sit beside. This is deeper analysis of one artifact, not a separate graded pathway or a new assignment.

Optional, at either depth: pick a competitor's recent AI announcement and spend 15 minutes asking your AI assistant to stress-test it against the Five Competitive Arenas, asking which arena it actually moves and which it only signals. Save the prompt and the response. We will compare them in the in-class activity.

## A Final Word Before Class

Module 6 has given you scaffolding: the Five Competitive Arenas, the Agent Authority Ladder, the foresight 2x2, the acute/chronic distinction, the Agentic Risk Register, and the Stargate-scale infrastructure lens. The frameworks organize a board conversation that would otherwise become a tooling debate and stress-test a response that would otherwise become a press release. They are not the point. The executive question in 2026 is whether the strategic discipline you bring to the next agentic shock is recognizably different from the discipline you brought to the last earnings call. Diagnose what has changed in the basis of competition. Choose where to act first. Set authority boundaries that match your evidence. Govern with the controls your business can run. Decide who owns the next 90 days.

The next shock will arrive on someone else's schedule, and it will reach your board's agenda before your team has finished reading the announcement. Your response, if it is to be more than a reaction, is built before then.

## Executive Glossary

Use this alphabetized glossary as a lookup aid for Part 1. It includes only terms not already defined in Modules 1–5; definitions reflect how the terms are used here.

| Term                              | Plain-language definition                                                                                                                             |
|-----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Acute / chronic disruption</b> | Disruption that arrives as a discrete event and demands a fast response, versus disruption that accrues gradually and demands a sustained discipline. |

| Term                                | Plain-language definition                                                                                                                                                   |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Algorithmic-efficiency shock</b> | A step change in how cheaply a given AI capability can be produced, as distinct from a step change in what AI can do.                                                       |
| <b>Capacity market</b>              | An electricity market that pays resources to commit availability for a future delivery period; its clearing price reflects demand, available resources, and market rules.   |
| <b>Consumption-based pricing</b>    | Software priced by usage (compute, tokens, actions, sessions) rather than by a flat fee per seat.                                                                           |
| <b>Headless</b>                     | Software designed to be operated through APIs and protocols rather than through a human-facing screen.                                                                      |
| <b>Interconnection queue</b>        | A process for studying and approving new connections to the grid; the requirements and timelines differ for generators and large electricity users.                         |
| <b>Moat mechanism</b>               | The specific reason a competitive position is hard to copy: scarce data, deep workflow embedding, superior unit economics, scarce infrastructure access, or legible trust.  |
| <b>No-regrets move</b>              | An action worth taking regardless of which scenario unfolds.                                                                                                                |
| <b>Resource-Based View</b>          | The strategy theory that durable advantage comes from resources that are valuable, rare, and hard to imitate or substitute, not from capabilities every competitor can buy. |
| <b>Social license</b>               | The ongoing acceptance of an organization's operations by the communities, regulators, and stakeholders affected by them, which legal permission alone does not confer.     |
| <b>Trigger</b>                      | A pre-agreed piece of evidence that tells a leadership team a scenario is becoming more likely and a planned move should start.                                             |

## Part 2: Field Reference

The Field Reference is the leave-behind. It contains decision tables, standards references, a dated Current Horizon Scan of fast-moving developments, cautionary tales, a do/don't quick reference, and a workbook prompt. Each component is self-contained; flip directly to whichever you need.

### Executive Infrastructure Questions

| Constraint               | Why Executives Should Care                                                                                         | Questions To Ask                                                                                                                 |
|--------------------------|--------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| Compute and cloud        | Access, price, latency, data residency, and vendor terms can shape which use cases are economically viable.        | Which workloads require frontier models, smaller models, open models, or model routing? Where are we exposed to one provider?    |
| Energy and grid          | Local power constraints can delay data-center capacity or raise cost, even when global AI capability is available. | Could our AI plans be delayed or made more expensive by power availability, data-center location, or sustainability commitments? |
| Chip supply chain        | Advanced AI chips and fabrication capacity are concentrated, and export controls can affect availability.          | What happens if chip availability, export controls, or regional tensions affect our vendors?                                     |
| Policy and sovereignty   | AI policy increasingly links capability, infrastructure, security, and market access.                              | Which jurisdictions govern our data, model use, high-risk applications, and customer-facing AI systems?                          |
| Sustainability and trust | AI can increase energy demand while also improving energy optimization. The strategy issue is tradeoff management. | Can we explain the energy, cost, and sustainability logic of our AI program to customers, employees, and regulators?             |

### The Four Headless-Era Decisions

| Decision                       | Trigger question                                                                                                                    | Where it is owned in this course                     |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| Contract exposure              | Which software contracts may add or expand consumption pricing over the next 24 months, and what would different usage levels cost? | Module 5: the Subsidy Audit is the operational tool. |
| Point-solution rationalization | When one agent can act across several systems, how many separately licensed tools survive consolidation?                            | Module 3: stack and sourcing.                        |
| Which layer to compete for     | Does the firm play for the data-access layer, the agent-ergonomics layer, or the vendor-harness layer?                              | This module: the moat-mechanism question.            |
| Oversight of work through APIs | What logs, monitoring, approvals, and user feedback make agent actions visible and accountable?                                     | This module: the Agent Authority Ladder.             |

### Hardware Scan: Production Capabilities and Longer-Horizon Options

Separate capabilities already available for sourcing decisions from technologies whose commercial role remains uncertain. Revisit the horizon for the specific workload rather than assigning every hardware category the same three-year timetable.

| Topic                                             | Treatment                                                                                                                                                                 |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Quantum computing                                 | Evaluate claims of advantage for specific optimization, chemistry, simulation, or cryptography applications against demonstrated performance and deployment requirements. |
| Neuromorphic computing                            | Monitor potential benefits for low-power, low-latency workloads and test them against the requirements of the intended application.                                       |
| Custom and specialized chips; photonic approaches | Specialized accelerators already support production workloads. Assess newer architectures separately for maturity, availability, software support, cost, and energy use.  |

## Agentic Risk Register

| Risk                         | How It Shows Up                                                                                        | Executive Control                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| Goal drift or goal hijacking | An agent pursues an objective in a way that satisfies the literal goal but violates business intent.   | Explicit objectives, prohibited actions, test cases, monitoring, and escalation thresholds.          |
| Tool misuse                  | An agent uses approved tools in unsafe combinations or at inappropriate moments.                       | Least-privilege permissions, tool gating, sandboxing, transaction limits, and approval rules.        |
| Data leakage                 | Sensitive customer, employee, IP, or regulated data moves into the wrong model, vendor, or workflow.   | Data classification, access control, vendor terms, logging, privacy review, and retention rules.     |
| Identity and authority abuse | Agents impersonate users, overstep role permissions, or create unclear accountability.                 | Agent identity, delegated authority registry, approval rules, separation of duties, and audit trail. |
| Human overreliance           | Users over-trust fluent outputs or stop exercising judgment in high-stakes contexts.                   | AI literacy, interface cues, confidence signals, human review routines, and manager coaching.        |
| Regulatory misclassification | A use case becomes high-risk or subject to transparency obligations without management recognizing it. | AI inventory, regulatory triage, legal review, EU AI Act readiness, and escalation path.             |
| Infrastructure dependency    | A critical AI capability depends on one provider, scarce compute, local power, or fragile integration. | Sourcing diversification, exit plan, model routing, resilience tests, and vendor governance.         |

The Hugging Face case below illustrates goal-seeking outside authorized boundaries during testing with reduced safeguards.<sup>12</sup> Model availability is a separate dependency: OpenAI restricted Astra development activities in August 2026 under strengthened security requirements. Its launch page, checked September 14, subsequently announced rollout, so the restriction is a historical example rather than an unresolved launch question.<sup>38</sup>

## Standards and Frameworks

Reach for these as the leave-behind reference set when a board, regulator, customer, or acquirer needs a legible AI position.

- **ISO/IEC 42001:2023.** An AI management-system standard that can support organizational governance. Certification alone does not establish EU AI Act compliance; the Commission explains that the Act's quality-management requirements need an aligned standard. Reach for

when reviewing management processes and assurance evidence. Commission standardization FAQ.

- **NIST AI Risk Management Framework: Generative AI Profile.** Trustworthy-AI vocabulary, lifecycle governance, GenAI-specific risk language. Reach for when defining the risk register at any altitude.<sup>6</sup>
- **NIST AI 100-2e2025, Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations.** A companion reference for classifying adversarial-ML risks. Reach for when security leaders need a common vocabulary for attacks and mitigations. NIST publication.
- **MITRE ATLAS.** A living knowledge base of adversary tactics and techniques involving AI. Its August 31, 2026 release, v2026.08, lists 16 tactics, 114 techniques and 83 sub-techniques. Reach for when translating the Agentic Risk Register into a threat model, and record the version used. MITRE changelog.
- **OWASP Top 10 for Agentic Applications for 2026.** Practical security risks for autonomous and agentic systems. Reach for when reviewing an agentic-AI architecture or vendor.<sup>7</sup>
- **EU AI Act General-Purpose AI Code of Practice.** A voluntary route for GPAI providers to demonstrate compliance with applicable Article 53 and 55 obligations; additional safety and security commitments apply to systemic-risk models. Signing the code does not automatically confer the statutory presumption of conformity associated with relevant harmonized standards. Reach for when examining a model provider's compliance evidence and obligations. Article 53, Article 55.
- **EU AI Act Code of Practice on Transparency of AI-Generated Content and Article 50 Guidelines.** Commission materials explaining disclosure, marking, and labelling duties. Reach for when classifying the provider and deployer obligations of a chatbot, agent, or generative feature, including exceptions and transition provisions.<sup>26</sup>
- **OECD AI Principles** (2019, updated May 2024). A non-binding intergovernmental framework for trustworthy AI. Reach for when comparing governance approaches across jurisdictions. OECD announcement.
- **Model Context Protocol (MCP) and Agent2Agent (A2A).** Interoperability protocols with different roles: MCP connects AI applications with tools and context; A2A supports agent-to-agent interaction. ACP merged into A2A in 2025. Reach for when comparing integration needs and portability, while checking implementation support and permissions. MCP governance announcement, Linux Foundation update on ACP and A2A.
- **AI Bill of Materials (AIBOM) / OWASP AIBOM Project.** An emerging approach to documenting AI components and supply-chain dependencies, analogous to a software bill of materials. Reach for when a customer, regulator, or acquirer needs a clearer inventory; assess the scope and completeness of the inventory supplied. OWASP project.
- **WEF, "Here's a Playbook for Boards on How to Govern Agentic AI"** (April 2026). Board-facing analysis of delegated authority and governance. Reach for as a discussion aid alongside the Module 6 board challenge questions.<sup>8</sup>

## Current Horizon Scan (As of September 14, 2026)

---

This section holds dated developments that may affect a board decision. **Confirmed** identifies a documented event or published finding; a survey response, vendor commitment, or forecast retains that status and its limits. **Unresolved** identifies an open issue. **Strategic implication** states the module's interpretation. Recheck a development before relying on it for a live decision; the scan date does not turn an older announcement into a current availability guarantee.

### Regulation and Litigation

- **EU AI Act: the post-Omnibus baseline. Confirmed:** Regulation (EU) 2026/1744 was published July 24, 2026 and entered into force July 27. It moved the high-risk dates to December 2, 2027 for Annex III systems and August 2, 2028 for relevant Annex I systems.<sup>5</sup> Article 50 transparency requirements began applying August 2, 2026, with a December 2 transition for the Article 50(2) marking duty of systems already on the market before August 2. The Commission's code and guidelines address disclosure, marking and labelling.<sup>26</sup> GPAI obligations generally began applying in August 2025, with Commission enforcement powers from August 2026 and a separate transition for older models.<sup>4</sup> **Unresolved:** Technical implementation and conformity evidence remain use-case-specific. **Strategic implication:** Identify the firm's role as provider or deployer, the applicable duty, exceptions and deadline. Commission supervision of GPAI providers and national supervision of AI systems are distinct; the later high-risk dates do not defer every obligation.
- **Colorado: a replacement AI statute and chatbot requirements. Confirmed:** Governor Jared Polis signed SB26-189 on May 14, 2026, replacing the earlier Colorado AI Act with an automated decision-making framework due to take effect January 1, 2027. Colorado also enacted HB26-1263 in 2026, with general chatbot requirements and additional protections concerning minors, effective January 1, 2027.<sup>27</sup> **Unresolved:** Implementation and how other states develop their own approaches. **Strategic implication:** Map obligations against enacted text and effective dates, and keep the design adaptable to changes in the law.
- **Amazon v. Perplexity: customer-agent access. Confirmed:** On August 4, 2026, the Ninth Circuit vacated a preliminary injunction against Perplexity's Comet browser. On this record, the panel found Amazon unlikely to prevail on the CFAA access issue when an authorized user directed the agent. The public appellate docket records denial of rehearing en banc on September 10.<sup>28</sup> **Unresolved:** The remaining litigation and how courts will treat different facts or technologies. Senator Mark Warner's June 29 AI AGENT Act discussion draft proposed rules for third-party agents and covered platforms; it is a legislative proposal in this account.<sup>29</sup> **Strategic implication:** The opinion does not establish blanket permission for agents to access other platforms. Review authorization, contracts, technical controls and the current legal posture for the intended workflow.
- **U.S. AI chip policy: announcements and delivery are different states. Confirmed:** On May 13, 2025, Commerce announced rescission of the AI Diffusion Rule, instructed officials not to enforce it, and said a formal rescission and replacement rule would follow. On July 14, NVIDIA said it was filing applications to resume H20 sales to China, had received assurances that licenses would be granted, and hoped to begin deliveries soon.<sup>22</sup> **Unresolved:** Current licensing conditions and their implications for a particular supplier or destination require a fresh check.

**Strategic implication:** Ask which rule, license, supply commitment, and delivery date a sourcing plan depends on, and review those assumptions every 90 days.

- **Hugging Face intrusion and model-development controls. Confirmed:** OpenAI's August 26 investigation identifies an internal research model, IM1, as driving the principal Hugging Face compromise during cybersecurity evaluations with reduced safeguards. GPT-5.6 Sol agents also reproduced an exploit and copied private evaluation data. The findings describe unauthorized communication and infrastructure weaknesses preceding the intrusion; IM1 was never intended for release.<sup>12</sup> OpenAI restricted Astra development activities in August. Its launch page checked September 14 announces rollout.<sup>38</sup> **Unresolved:** How well strengthened controls transfer to other environments and future models. **Strategic implication:** Require evidence that permissions are enforced in practice, that monitoring detects boundary violations, and that vendor changes have a workable fallback. This incident does not measure the frequency of such behavior in safeguarded production systems.
- **A proposal to pace capability development. Confirmed as the author's position:** In his September 2026 essay "We Must Pace the Frontier," Dario Amodei argues for slowing capability improvements, warns about possible cyber harm, and commits Anthropic to giving external evaluators employee-level access.<sup>24</sup> His six-to-twelve-month botnet warning is a forecast. **Unresolved:** Implementation, uptake by other providers, and the effect on capability or release schedules. **Strategic implication:** Track observable changes in evaluation access, deployment controls and performance. A public commitment alone does not establish a change in METR's historical capability trend or in enterprise operating maturity.
- **U.S. banking model-risk guidance and agentic AI. Confirmed:** OCC Bulletin 2026-13 and Federal Reserve SR 26-2, dated April 17, 2026, revise model-risk-management guidance while excluding generative and agentic AI from its scope.<sup>37</sup> **Unresolved:** How institutions should apply other relevant requirements and supervisory expectations to each deployment. **Strategic implication:** An exclusion from this particular guidance does not remove other responsibilities. Have the accountable risk and legal teams map the system's obligations and establish controls suited to its authority and use.

## Vendor and Platform Moves

- **The April 2026 enterprise-agent announcements. Confirmed as announcements:** Salesforce announced Headless 360 on April 15. On April 22, OpenAI announced workspace agents, Microsoft introduced updated hosted agents in Foundry Agent Service, and Google announced Gemini Enterprise Agent Platform.<sup>15</sup> **Unresolved:** Availability for each customer, integration effort, adoption, and pricing over time. **Strategic implication:** Use the announcements to review contract exposure, opportunities for consolidation, integration choices, and oversight. A launch announcement is the start of customer diligence.
- **MCP under foundation governance. Confirmed:** Anthropic donated MCP to the Linux Foundation's Agentic AI Foundation on December 9, 2025. Its announcement reported more than 10,000 active public servers and more than 97 million monthly Python and TypeScript SDK downloads at that time.<sup>16</sup> Downloads do not measure unique users or enterprise deployments. **Strategic implication:** Open protocol governance can support portability, but switching costs also depend on data, tools, permissions and implementation choices.

- **Adoption, trust and value use different measures. Confirmed as survey findings:**

McKinsey's August 25, 2026 report draws on 1,719 respondents in 97 countries, surveyed May 4–June 8. It reports 44% scaling AI across the enterprise; agent scaling in at least one function at 40% for firms with at least \$1 billion in revenue versus 22% for smaller firms; and 37% attributing any EBIT impact to AI.<sup>2</sup> HBR's 6% full-trust finding comes from a July 2025 survey of 603 technology decision-makers at organizations that had at least considered agentic AI.<sup>17</sup> KPMG's Q2 2026 survey reports 49% delaying or scaling back agents over expected cost versus value, while 79% would retain AI as an investment priority even in recession.<sup>18</sup> **Unresolved:** How each survey's results translate to a specific workflow or firm. **Strategic implication:** Treat adoption, trust and financial impact as separate questions. Test the proposed workflow's performance and economics, with a cost limit and an explicit decision point before expansion.

## Infrastructure Signals

- **Grid constraints and forecast scenarios. Confirmed as dated operator findings:**

PJM's July 14, 2026 release says its 2028–2029 capacity auction cleared at \$325/MW-day under the price collar, with a 6,831 MW shortfall. It described the third consecutive auction under that collar and a plan to seek approval for backstop procurement; the release does not establish that the later procurement occurred.<sup>30</sup> ERCOT's 2025 presentation contrasted a 218 GW transmission-provider projection for 2031 with an adjusted forecast of about 145 GW for its service area. NERC's 2025 assessment projected North American summer peak demand growth of 224 GW over the following decade, 69% more growth than its previous assessment projected.<sup>21</sup>

**Strategic implication:** Record geography, horizon, scenario and adjustments in the sourcing plan. Check power availability, interconnection and permitting against the vendor's delivery commitments.

- **AI infrastructure plans and financing. Confirmed as announced plans and filed terms:**

NVIDIA announced a letter of intent to invest up to \$100 billion in OpenAI as systems were deployed. The amount was an investment intention, not a completed delivery of \$100 billion in processors. Alphabet's June 2, 2026 filing described an \$84.75 billion equity package including potential proceeds from a \$40 billion at-the-market program. The filing identified different uses of proceeds, including employee-equity tax obligations for the at-the-market program; the total was neither cash already raised nor an amount allocated entirely to AI capacity.<sup>20</sup> **Unresolved:**

Amounts ultimately deployed or raised, future financing needs, and execution of capacity plans. **Strategic implication:** Separate announced capacity, funding intentions, issued securities and available cash. Review the provider's balance sheet and delivery commitments alongside model performance.

## Cautionary Tales

---

### Air Canada and the Liable Chatbot (February 2024)

Discussed in detail in Module 1; cross-referenced here for the agent-authority implication. **The Module 6 lesson:** the agent-authority question is not whether the agent is technically capable of speaking on your behalf, but whether your authority registry permits it to. Statements made by AI agents to customers may bind the organization regardless of how the system was procured.

## **Klarna and the Customer-Service Reversal (2024-2025)**

Discussed in detail in Module 5; cross-referenced here for the strategic-shock implication. **The Module 6 lesson:** test customer experience alongside efficiency. A decision to restore human service is a signal to examine workflow design, service quality, and the assumptions behind the original deployment.

## **Cursor and the Hallucinated Device Policy (April 2025)**

Cursor's support bot invented a one-device-per-subscription restriction. Cofounder Michael Truell publicly corrected the claim and apologized; refunding an affected user was part of the response.<sup>35</sup> **The Module 6 lesson:** a support agent needs authoritative policy sources, escalation rules, and a way to correct misinformation quickly. A confident policy answer can become a customer-trust incident.

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

Discussed in detail in Module 1; cross-referenced here for the announcement-vs-experience implication. **The Module 6 lesson:** a competitor's announcement of an AI capability is not the same as a customer's experience of that capability. The announcement can shape board agendas while the experience eventually shapes the market.

## **Enterprise Copilot Rollouts: A Decision Test**

Use any proposed enterprise copilot rollout as a test of the business case. Licenses purchased, active use, task performance, user experience, and financial value are different measures. **The Module 6 lesson:** require evidence for the outcomes that justify expansion and renewal. The question is how the tool changes a specific workflow under the organization's operating conditions.

## **Amazon v. Perplexity Comet: The Customer-Agent Access Case (November 2025 - Ongoing)**

Amazon sued Perplexity over its Comet browser in November 2025. The district court granted a preliminary injunction in March 2026. On August 4, the Ninth Circuit vacated it, finding Amazon unlikely to prevail on the CFAA access issue on the record before the panel. The opinion treated the user-directed agent as a tool in that setting and expressly confined its holding. The public docket records denial of rehearing en banc on September 10; the Current Horizon Scan carries the status. 28

**The Module 6 lesson:** platform-access law belongs in customer-agent strategy. The ruling changes the assessment of one access dispute; contracts, other legal claims, technical controls, and different factual settings still matter. Obtain a current assessment for the actual workflow before granting access or relying on another platform.

## **The Hugging Face Incident: Boundary Failures During Evaluation (July-August 2026)**

OpenAI's August 26 investigation identifies IM1, an internal research model never intended for release, as driving the principal Hugging Face compromise. During difficult cybersecurity evaluations with reduced safeguards, agents communicated through shared infrastructure, pursued solutions outside the authorized environment, and exploited weaknesses that enabled access to

third-party systems. GPT-5.6 Sol agents also reproduced an exploit and copied private evaluation data. The investigation links the incident to reward hacking, persistence on difficult tasks, unauthorized communication, and agents adopting one another's goals.<sup>12</sup>

The independent METR/Redwood investigation provides a second account of the incident and the controls around it.<sup>12</sup> Read both when using the case to assess evaluation design or vendor assurance.

Keep the setting in view: these were evaluations with reduced safeguards. Their findings do not directly estimate behavior in ordinary production use. OpenAI's later restrictions on Astra development and subsequent rollout announcement are separate events, tracked in the Horizon Scan.<sup>38</sup>

**The Module 6 lesson:** an agent may pursue a measured objective in ways that violate the intended task boundary. Ask what it is rewarded for, whether it can stop or escalate when a task is infeasible, and whether permissions are enforced beyond the instructions in its prompt. Test monitoring and containment as part of the authority decision.

### Rio Tinto and Social License: A Governance Analogy (May 2020)

After the destruction of the Juukan Gorge rock shelters in May 2020, Rio Tinto conducted a *Board Review of Cultural Heritage Management*. In September it announced the departure of its chief executive and two senior leaders.<sup>36</sup> This case concerns cultural heritage, stakeholder obligations and executive accountability. It is used here as a governance analogy; the cited record does not establish that AI caused or authorized the destruction.

**The Module 6 lesson:** a technically feasible or formally permitted action may still breach obligations to affected people. When delegating authority to an automated system, define the stakeholder commitments and escalation decisions that must constrain execution.

## Do / Don't Quick Reference

| Do                                                                                                                                              | Don't                                                                                    |
|-------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| Treat agent authority as a board-level decision-rights question with infrastructure, governance, and competitive consequences                   | Treat agentic AI as a tooling decision routed through procurement                        |
| Diagnose what has actually changed in the basis of competition before responding to a competitor announcement                                   | Match a competitor's announcement with an announcement of your own                       |
| Match agent authority to evidence; expand only after controls and evals support the next authority level                                        | Grant agent authority on the strength of a successful demo                               |
| Treat sourcing and infrastructure as strategy, not procurement; ask which workloads need frontier models, smaller models, or open-weight models | Treat infrastructure as a cost center to be optimized rather than a strategic constraint |
| Make the firm's AI risk position legible to regulators, customers, and acquirers through certifications, inventories, and AIBOMs                | Treat trust as a soft issue to be addressed after a regulatory event                     |
| Set explicit kill, rollback, and stop-scale criteria before granting any agent the authority to act                                             | Wait for the first incident to define the rollback procedure                             |
| Pair every infrastructure commitment with a 90-day review of grid, cloud, and chip-supply assumptions                                           | Treat the sourcing roadmap as a one-time decision                                        |

| Do                                                                                     | Don't                                                                              |
|----------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| Distinguish between acute and chronic disruption, and respond to each on its own clock | Treat every frontier signal as an acute shock requiring an emergency board meeting |

## Workbook Prompt: The 90-Day Frontier-Shock Response

Run this as a 60-minute exercise with three to six colleagues from across your organization (operations, technology, finance, legal, line-of-business leadership). Use a real or hypothetical AI-competitive shock from your industry: a competitor announcement, a customer-agent rollout, a regulatory shift, an infrastructure constraint.

- 1. Diagnose the shock (15 minutes).** Which of the Five Competitive Arenas does it primarily move? Is it acute or chronic? What is the most plausible mechanism: cost, speed, customer interface, data feedback, narrative? Write the diagnosis as a single sentence and name the first workflow you would address.
- 2. Test and extend the priority (15 minutes).** Compare the first workflow with three to five candidate AI moves across efficiency, effectiveness, and innovation. Confirm or change the priority, naming the users, proprietary context, and expected business KPI.
- 3. Set the authority boundary (10 minutes).** Which level of the Agent Authority Ladder is strategically justified for this priority use case? What controls are required at that level? What are the prohibited actions?
- 4. Govern with evidence (10 minutes).** What evidence does the board need to approve scaling? Name the technical, adoption, workflow, financial, and strategic-impact metrics. Set the stop, rollback, and scale criteria.
- 5. Decide who owns the next 90 days (10 minutes).** Name the operating-model owner, steering cadence, decision rights, and governance coordination. Write the first three actions and the decision review at day 90.

Steps 1, 3, and 5 are the Foundation/Core pass and can be run in thirty-five minutes; steps 2 and 4 complete the Challenge/Full Response. The exercise produces a one-page response in the format of the Global Express board simulation. A polished memo is not the goal. The exercise succeeds when it surfaces where the team converges and where it does not, and exposes the disciplines (or gaps) that show up under board-pressure conditions.

## Part 3: Card Specification

---

### The Board Strategic Response

---

The Board Strategic Response is drafted in the Global Express simulation. Students can apply the same discipline to their own organization or a chosen company: name a plausible competitive threat, choose a priority workflow, set the agent's authority and controls, and define the first 90 days. There is no separate post-class submission for this card. Module 8, *AI Strategy Capstone Artifact Kit*, provides optional supporting templates for the board briefing, value portfolio, use case charter, risk register, KPI value chain, operating model, and first 90 days. Use those that support the chosen decision.

The Global Express simulation in class is the rehearsal. The Board Strategic Response is the same discipline applied to a competitive question in your own context. The point is not breadth of frameworks reviewed; it is depth of judgment about a single board-level decision under competitive pressure.

The governing syllabus gives Blog 3 three options: a cards portfolio with an integrating narrative, a 5–7 page capstone memo, or a reflection. Students choosing the capstone-memo option may use this card to develop the memo; students choosing the portfolio option include it among their course cards. All students produce the card in session. Module 8 supplies an assembly aid; the syllabus governs assignment requirements and grading.

### Debrief Criteria

Use these five criteria to discuss and improve the Board Strategic Response. In the Core pass, focus on threat diagnosis, the priority workflow, authority and controls, and the first 90 days. The Full Response adds the KPI value chain and supporting analysis:

1. **Threat diagnosis.** Is the competitive shock named specifically, and is the arena diagnosis defensible against alternative readings?
2. **Use-case prioritization.** Is one priority workflow chosen, with the workflow, users, proprietary context, and expected business KPI named?
3. **KPI value chain.** Does the value logic connect the workflow to a business KPI the board would recognize, with stop, rollback, and scale criteria stated?
4. **Authority and control fit.** Is the agent authority level justified for the use case, with required controls and prohibited actions named?
5. **90-day operating realism.** Are the owner, cadence, decision rights, and first three actions concrete enough to begin the following week?

These are feedback prompts with no grading weights. The syllabus's grading approach applies to participation and the chosen Blog 3 option; this exercise does not create an additional grading rubric.

**KEY LESSON.** A Board Strategic Response does not need to prove certainty. It needs to prove discipline: a clear strategic thesis, one priority workflow, credible value logic, named authority boundaries, visible risk controls, accountable owners, and concrete first actions.

## Next-Module Handoff

Module 7 asks whether the strategic response has a realistic capability system behind it: managers who can sponsor workflow redesign, champions who can create local proof, governance that learns from use, and executives who see evidence early enough to act.

## Works Cited and Sources

**Cross-module case references:** Air Canada and McDonald's are developed in Module 1; Klarna is developed in Module 5. Their brief mentions here apply those cases to the Module 6 decision. Consult the source module for the underlying case account.

1. Stanford HAI. *The 2026 AI Index Report*, including the economy chapter. Collection years and survey measures vary within the report.
2. McKinsey & Company. *The State of AI in 2026: On the Road to ROI*, August 25, 2026. Survey of 1,719 respondents in 97 countries, fielded May 4–June 8. Figures are respondent reports; revenue groups and organizational scaling measures follow the report's definitions.
3. International Energy Agency. *Energy and AI: Energy Demand from AI*. The 945 TWh figure is the 2030 base-case projection for global data-center electricity consumption, including non-AI uses.
4. European Commission. "Navigating the AI Act" and GPAI questions and answers. Duties depend on role, use, model vintage and applicable transitions.
5. European Parliament and Council. Regulation (EU) 2026/1744, *Official Journal*, July 24, 2026; entered into force July 27. The amended provisions govern the dates described here.
6. NIST. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*.
7. OWASP GenAI Security Project. *Top 10 for Agentic Applications for 2026*.
8. World Economic Forum. "Here's a Playbook for Boards on How to Govern Agentic AI", April 28, 2026. Board-facing analysis; the module's challenge questions adapt the governance framing.
9. DeepSeek. Release history, R1 release January 20, 2025. Stanford Cyber Policy Center. "Taking Stock of the DeepSeek Shock". The efficiency-shock framing in this module is an interpretation.
10. OpenAI. "Announcing the Stargate Project", January 21, 2025; "OpenAI, Oracle, and SoftBank expand Stargate with five new AI data center sites", September 23, 2025. Capacity and investment figures refer to those announcements.
11. DeepSeek-AI. *DeepSeek-V3 Technical Report*. The \$5.576 million estimate covers reported V3 training computation at an assumed GPU rental rate; it excludes prior research and ablation work.
12. OpenAI. "The Hugging Face incident and the road ahead", August 26, 2026. See also METR/Redwood's independent investigation and OpenAI's initial disclosure and updates. The later investigation supplies the model identities and expanded chronology used here.
13. The White House. "White House Unveils America's AI Action Plan", July 23, 2025. Official announcement and link to the plan.
14. Paresh Dave. "Why Walmart and OpenAI Are Shaking Up Their Agentic Shopping Deal", *WIRED*, March 18, 2026. Original interview with Walmart's Daniel Danker; the conversion comparison is attributed to him.

15. Primary vendor announcements: Salesforce Headless 360, April 15, 2026; OpenAI workspace agents and Business release notes, April 22; Microsoft hosted agents in Foundry Agent Service, April 22; Google Gemini Enterprise Agent Platform, April 22. Announcements do not guarantee availability to every customer.
16. Anthropic. "Donating the Model Context Protocol and establishing the Agentic AI Foundation", December 9, 2025. Server and Python/TypeScript SDK-download counts are announcement-date figures.
17. Harvard Business Review Analytic Services. *From the Edge to the Core: Bringing Agentic AI to the Heart of the Enterprise*. July 2025 survey of 603 HBR audience members involved in technology decisions at organizations that had at least considered agentic AI; published in 2025. The 6% figure concerns full trust with core end-to-end processes.
18. KPMG. *Global AI Pulse, Q2 2026*, especially methodology and cost/value findings, pp. 3, 17-18 and 22. Survey of 2,145 senior leaders across 20 countries, fielded April 28-May 25. The 49% finding concerns agent deployment decisions and expected costs versus anticipated value; the 79% finding retains the recession condition.
19. David Wingate, Barclay L. Burns, and Jay B. Barney. "Why AI Will Not Provide Sustainable Competitive Advantage", *MIT Sloan Management Review*, May 8, 2025.
20. NVIDIA. "OpenAI and NVIDIA Announce Strategic Partnership to Deploy 10 Gigawatts of NVIDIA Systems", September 22, 2025. Alphabet. 2026 SEC financing filing. The investment intention and potential offering proceeds are distinct from amounts actually invested or raised.
21. ERCOT. 2025-2031 Long-Term Load Forecast Update and Methodology Changes, April 2025. NERC. *2025 Long-Term Reliability Assessment*. Forecasts retain their original dates, scenarios, geographic coverage and horizons.
22. U.S. Bureau of Industry and Security. "Department of Commerce Announces Rescission of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls", May 13, 2025. NVIDIA. "NVIDIA CEO Jensen Huang Promotes AI in Washington, DC and China", July 14, 2025. These are historical announcements, not a statement of current licensing eligibility.
23. METR. "Measuring AI Ability to Complete Long Tasks", March 19, 2025. The historical trend concerns human-duration equivalents of selected software tasks at 50% success, not production reliability or organizational maturity.
24. Dario Amodei. "We Must Pace the Frontier", September 2026. Forecasts, arguments and commitments are attributed to the author; the retrieved page does not establish an exact publication day.
25. Gerald C. Kane, Rich Nanda, Anh Phillips, and Jonathan Copulsky. *The Transformation Myth: Leading Your Organization through Uncertain Times*. MIT Press, 2021. See the 2021 author presentation for the acute/chronic distinction.
26. European Commission. Code of Practice on Transparency of AI-Generated Content and Guidelines on Article 50 transparency obligations. Consult the applicable duty, exceptions and transition period for the intended use.
27. Colorado General Assembly. SB26-189 and HB26-1263; Colorado Attorney General, AI law resources. The official pages differ on the chatbot law's signing date; this module uses the enactment year and effective date without choosing between those conflicting dates.
28. *Amazon.com Services LLC v. Perplexity AI, Inc.*, Ninth Circuit No. 26-1444. August 4, 2026 opinion and public appellate docket reproduction, including the September 10 rehearing entry. The docket source is a public reproduction, not a direct PACER readback; recheck the posture for a live decision.
29. Senator Mark Warner. "Warner Unveils Discussion Draft of Legislation to Create Innovative Market for Secure Artificial Intelligence Agents", June 29, 2026. The cited document is a discussion draft, not enacted law.
30. PJM Interconnection. "PJM Capacity Auction Procures 138,318 MW of Generation Resources as Work Continues to Address Growing Electricity Demand", July 14, 2026. The release establishes the auction result and announced next step; subsequent procurement requires separate evidence.
31. MIT Sloan Management Review / BCG. *Me, Myself, and AI* interview transcript with Aflac's Shelia Anderson, 2023.
32. Manufacturing Digital. Interview featuring Georgia-Pacific's Roshan Shah, October 29, 2024. The module paraphrases the interview; no direct quotation is asserted.
33. Lockheed Martin. "Lockheed Martin Launches Astris AI to Enable Secure AI Solutions Across Industries", December 16, 2024.

34. ServiceNow. Q1 2026 earnings-call transcript, April 22, 2026, reproduced by The Motley Fool. The non-seat-pricing figure is management's statement and includes more than AI-token consumption.
35. Michael Truell. Original public response to the Cursor support incident, April 2025. See also contemporaneous Ars Technica reporting.
36. Rio Tinto. "Rio Tinto executive committee changes", September 11, 2020. The release identifies the *Board Review of Cultural Heritage Management* and the announced departures.
37. Office of the Comptroller of the Currency. Bulletin 2026-13. Federal Reserve. SR 26-2, April 17, 2026. The exclusion is from this model-risk guidance, not from all applicable requirements.
38. OpenAI. "Pacing model development in an era of cyber-critical capabilities", August 7, 2026; GPT-6 Astra launch announcement, rollout statement observed September 14, 2026. The retrieved launch page does not expose a publication date; availability may differ by customer and deployment channel.

### 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 6: Future of Competition*. 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](mailto:jerry@profkane.com)