INSTITUTIONAL AI

AI CONTROL. FOR INSTITUTIONS.

The 25 Questions

Every Board Must Ask on AI Control

One question per cell of the 5×5 Control Matrix™ — the technical-evidence layer beneath the board's fiduciary duty.

BOARD BRIEFING | **2026 EDITION**

Companion to The State of AI Control in Institutional Finance, 2026 Edition

www.institutionalai.net

A note to the board

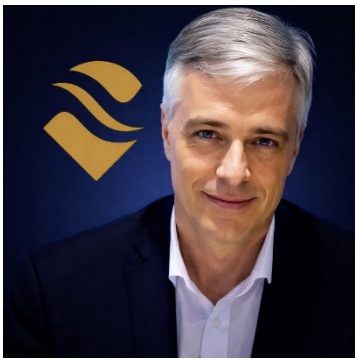

By Rad H. Pasovschi, Founder & CEO, Institutional AI, LLC

Governance describes what an institution intends to do. Control describes what it can demonstrably do. The distance between the two is where fiduciary and regulatory exposure accumulates. These twenty-five questions are designed to close that distance — not with policy attestations, but with evidence.

Each question corresponds to one cell of the 5×5 Control Matrix™ — five AI ecosystems (Power, Compute, Data Centers, Models, Agents) measured against five pillars of control (Jurisdictional, Logical, Technical, Operational, Contractual). Twenty-five cells, twenty-five questions. The board does not need to resolve the technical implementation of any cell; it needs management to answer the question with evidence rather than assurance.

The test is simple. For each question, management should be able to produce technical or contractual evidence on demand. Where the honest answer is "we rely on the provider's representations," the cell is governed by assurance, not controlled by the institution — and the board should record that as an open item with an owner and a date.

Scoring each answer

Controlled — evidence can be produced on demand, today.

Partial — evidence exists but requires manual reconstruction or covers only part of the estate.

Assured only — the institution relies on provider representations it cannot independently verify.

The 5×5 Control Matrix™ at a Glance

Five ecosystems, five pillars of control, twenty-five cells. The matrix below shows every cell in the document with its specification number. Each cell receives its own question for the board on the pages that follow.

Each cell begins with the single question whose answer determines whether the institution has technical control of the cell. The question is designed to be uncomfortable. If the question is uncomfortable to ask, the institution probably does not have control. If the question is comfortable to ask but the answer is uncertain, the institution definitely does not have control.

Pillars → ↓ Ecosystems	Jurisdictional	Logical	Technical	Operational	Contractual
Power	1.1	1.2	1.3	1.4	1.5
Compute	2.1	2.2	2.3	2.4	2.5
Data Centers	3.1	3.2	3.3	3.4	3.5
Models	4.1	4.2	4.3	4.4	4.5
Agents	5.1	5.2	5.3	5.4	5.5

> *Data is the connective tissue across all five ecosystems — it is the substance that flows through them, not a sixth ecosystem. The Five Pillars apply to data implicitly across every cell.*

ECOSYSTEM 1 — POWER

The electricity, grid, and carbon foundation beneath every AI workload.

1.1 Jurisdictional Control

Can we prove which grid, under which sovereign authority, is powering every AI workload at any given moment — and halt execution when that condition fails?

1.2 Logical Control

Do we know who can issue power-state commands — power-on, power-off, capacity allocation, emergency override — against our AI infrastructure, and is every such command authorized, logged, and reversible by us?

1.3 Technical Control

Can we cryptographically attest the generation source, carbon intensity, and grid jurisdiction of every kilowatt-hour our AI consumes — at hourly granularity, for any historical window?

1.4 Operational Control

Do we have real-time, workload-level visibility into the power consumption, carbon intensity, and grid-resilience profile of every AI deployment — with automated alerting when any exceeds sanctioned bounds?

1.5 Contractual Control

Do our contracts with facility, utility, and cloud providers grant enforceable rights to audit, attest, and challenge the power sourcing and carbon attribution of every workload — with material breach remedies if those rights are obstructed?

ECOSYSTEM 2 — COMPUTE

The silicon and processing substrate that executes AI workloads.

2.1 Jurisdictional Control

Can we prove, in real time and for any historical inference, the physical location and legal jurisdiction of the silicon that executed the workload — and prevent execution outside the sanctioned envelope?

2.2 Logical Control

For every privileged action against our AI compute — workload start, scaling, configuration, monitoring, debugging — can we produce the authenticated actor, the specific authorization, and the immutable record of what was done?

2.3 Technical Control

Can we prove every AI workload executes on verified, unmodified silicon, inside a confidential-computing boundary the provider cannot inspect, with the cryptographic trust root terminating with us rather than the provider?

2.4 Operational Control

Do we have continuous, immutable visibility into every aspect of our AI compute consumption — what executed, where, on what silicon, against what data, with what outputs and downstream actions?

2.5 Contractual Control

Do our compute contracts grant unconditional audit rights, enforceable exit provisions, subprocessor approval workflows, and breach remedies tied to specific operational conditions — not to "reasonable cooperation" language?

ECOSYSTEM 3 — DATA CENTERS

The physical facilities and hardware supply chain hosting the workloads.

3.1 Jurisdictional Control

Can we prove, with technical and legal evidence, the specific physical facility executing each workload, its jurisdictional designation, and the legal regime governing it — and prevent movement outside the sanctioned envelope?

3.2 Logical Control

For every individual with physical or privileged logical access to the facilities executing our AI workloads, can we produce the authenticated identity, the authorization scope, and the immutable audit trail of every access event?

3.3 Technical Control

Can we prove the integrity of the hardware supply chain that produced our hosting equipment, its tamper-resistance in operation, and the cryptographic identity of every device in the network path to our inference?

3.4 Operational Control

Do we have continuous visibility into the operational state of every facility hosting our AI — environmental conditions, equipment lifecycle, incidents, capacity, maintenance — at the granularity required for resilience decisions?

3.5 Contractual Control

Do our facility contracts grant enforceable rights to physical audit, sub-operator restriction, transfer-of-control protection, and media-destruction verification — with breach remedies tied to specific facility-layer conditions?

ECOSYSTEM 4 — MODELS

The AI models — proprietary and commercial — the institution invokes.

4.1 Jurisdictional Control

For every model we invoke, can we prove the jurisdiction of the provider, the jurisdictions of training and inference, the legal regime governing the training data, and the jurisdictional disposition of every prompt and output?

4.2 Logical Control

For every action that touches our models — deployment, invocation, configuration, fine-tuning, retirement — can we produce the authenticated identity, the specific authorization, and the immutable record of who did what, to which model, and when?

4.3 Technical Control

Can anyone other than us decrypt, extract, or run inference against our proprietary model weights? If yes — even theoretically, even by the provider, even under legal compulsion — we have governance language, not technical control.

4.4 Operational Control

For every inference from every model we operate, can we reconstruct on demand — for any historical window — the prompt, the data context, the output, the human review or override, and the downstream institutional action that followed?

4.5 Contractual Control

Do our model-provider contracts grant enforceable rights on inference-data use, model-change notification, output ownership, training-data exclusion, audit, and termination — with breach remedies appropriate to our exposure?

ECOSYSTEM 5 — AGENTS

Autonomous AI agents that take actions across institutional systems.

5.1 Jurisdictional Control

For every action an institutional agent takes, can we prove the jurisdiction in which the action occurred, the jurisdiction of the systems it touched, and the legal regime under which the agent itself operated?

5.2 Logical Control

For every agent we operate, can we define and enforce — with fine-grained, time-bounded, revocable authority — which actions it may take, against which systems, against which data, on whose behalf, and for how long?

5.3 Technical Control

Can we prove the technical integrity of every agent deployment — the attested execution environment, cryptographically signed action records, verified kill-switch capability, and rollback infrastructure — across every agent we run?

5.4 Operational Control

For every agent action, can we reconstruct on demand the full action context — the human principal, the systems touched, the data accessed, the outputs produced, the downstream effects, and any institutional review or override?

5.5 Contractual Control

Do our contracts with agent providers and the systems our agents invoke clearly allocate liability, indemnification, audit rights, and termination rights — with an architecture appropriate to the consequence of agent-action failure?

What the board should do with the answers

A board that asks these twenty-five questions and records honest answers will produce, in a single sitting, the most accurate picture of its institution’s real AI control posture it has ever held.

The pattern of answers — not any single answer — is the finding.

- **Identify the critical cells.** Five cells — one per ecosystem — typically carry the highest fiduciary consequence. Resolve those first.
- **Assign an owner and a date to every "assured only" answer.** An open control gap without an owner is an unmanaged risk.
- **Make this a standing review.** AI control is a fiduciary matter that sits above the technology-risk layer. The twenty-five questions are a quarterly instrument, not a one-time exercise.
- **Our recommendation.** This is not a short-duration engagement. The institutional commitment to a complete twenty-five-cell architecture is, in most cases, a multi-year program. Our firm's view is that this is appropriate. Institutional AI control is not a problem to be solved once and forgotten; it is a posture that must be maintained, audited, and evolved as the underlying technology, regulatory environment, and threat landscape change. The institutions that recognize this — and that build the institutional capacity to sustain the architecture as a permanent operational discipline — will be the institutions whose AI control posture remains defensible in 2027, 2028, and beyond. The institutions that treat AI control as a one-time project will find themselves repeating the exercise every eighteen months as the gap reopens.

About Institutional AI

The firm

Institutional AI is the AI control firm for institutional finance. We work with organizations where trust, accountability, and sovereignty are non-negotiable — asset owners, asset managers, asset servicers, banks, insurers, wealth managers, family offices, and private equity firms. AI control is a technical reality, not a policy exercise: governance is policy; control is technical reality.

We are unconflicted — by structure, not by assertion. We maintain no hyperscaler alliance, no model-vendor partnership, no resale economics with any infrastructure provider, and no commercial arrangement with any firm whose offerings we evaluate. That independence defines what institutional buyers can expect from our recommendations: counsel that answers to the institution, and to no one else.

Begin

Two steps, in either order.

Read the full analysis. *The State of AI Control in Institutional Finance, 2026 Edition* — the complete sector-by-sector assessment this briefing accompanies — is available on request.

Assess your own posture. The AI Sovereignty Assessment™ scores your institution against all twenty-five cells of the 5×5 Control Matrix™ and produces a board-ready picture of where asserted control and technical reality diverge.

Contact

Institutional AI, LLC · www.institutionalai.net · information@institutionalai.net

Disclaimer

This document is provided solely for informational and strategic discussion purposes and reflects Institutional AI's analytical opinions, framework interpretations, and operational perspectives as of the date of publication. It is not intended as legal, regulatory, investment, accounting, cybersecurity, or other professional advice.

The frameworks, classifications, maturity observations, and architectural concepts described herein are illustrative analytical constructs designed to support institutional evaluation and discussion. They are not certifications, guarantees, ratings, or representations of regulatory compliance, operational resilience, fiduciary sufficiency, or technical security.

Institutional conditions, legal obligations, infrastructure environments, and regulatory expectations vary materially across organizations and jurisdictions. Institutions should independently evaluate all architectural, governance, operational, and contractual decisions with appropriate professional advisors.