

THE AI SOVEREIGNTY ASSESSMENT

WEALTH MANAGEMENT EDITION (PARTIAL)

*Governing AI Sovereignty Across Client Advisory, Financial Planning,
Suitability, and the Client Relationship*

A Diagnostic Framework for Wealth Management Leadership

FOR THE BOARD | EXECUTIVE BRIEF | WEALTH MANAGEMENT | CONFIDENTIAL

This assessment maps your institution's AI sovereignty across 25 specific governance questions — the intersections of five control pillars and five AI infrastructure layers. For wealth managers, the stakes are deeply personal: the AI systems advising your clients have access to the most sensitive financial information those clients will ever share with any institution.

The governance architecture

INSTITUTIONAL AI AI SOVEREIGNTY ASSESSMENT THE 5x5 CONTROL MATRIX · 25 GOVERNANCE INTERSECTIONS · MAX SCORE 100 · © 2026 INSTITUTIONAL AI					
FIVE PILLARS × five ecosystems	1 — POWER Energy infrastructure	2 — COMPUTE GPU / chip infrastructure	3 — DATA CENTERS Cloud & physical infrastructure	4 — MODELS LLMs & AI systems	5 — AGENTS Agentic applications
PILLAR 1 Jurisdictional Control Where does it execute, under which law?	Where is energy data processed and under which jurisdiction? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Under which jurisdiction does GPU compute actually execute? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Can you prove where every data center workload executes? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Which jurisdiction governs model training, storage, and serving — and does it grant access to weights? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Where do agent actions execute and where do decision logs reside? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 2 Logical Control Who can access it, when, with what proof?	Who has privileged access to energy systems — and is it logged in your SIEM? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Who accesses GPU clusters — including provider support engineers? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Can you prove no unauthorized access in the past 18 months? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Can your model provider access your queries and outputs — in your logs? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Who can approve or halt agents — are all actions logged in your systems? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 3 Technical Control Who holds the encryption keys?	Do you control encryption keys for energy and ESG data? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Who controls keys for AI training data and model weights? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do you hold HYOK for all sensitive data — or does your cloud provider? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Does your model provider process your data in plaintext on their infra? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do you have cryptographic controls over what agents can access and output? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 4 Operational Control Do you have real-time visibility?	Real-time visibility into energy use and carbon intensity per workload? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Real-time visibility into compute utilization and cost across all GPUs? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Can you detect a residency violation or breach within minutes, not hours? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Real-time visibility into model behavior, quality, and decision provenance? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Can you monitor, pause, or audit every agent action in real time? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 5 Contractual Control Do you have enforceable rights?	Do energy contracts include audit rights, portability, and exit provisions? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do compute agreements protect against unilateral capacity restrictions? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do data center agreements include audit rights and deletion certification? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Do you own your query logs and outputs — or does the provider retain? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Do agent agreements cover liability for autonomous decisions and data exit? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4

Critical cell — universally low governance; most institutions at Score 1 (Reactive) → 4 (Sovereign) per cell · Max total: 100 pts

Score each cell 1 (Reactive) to 4 (Sovereign). The distribution reveals where governance is structurally sound and where client data is exposed — particularly in the Models and Agents columns where advisory AI operates.

The wealth management client data imperative

Your clients share their most sensitive information with you — net worth, tax position, estate structure, family dynamics, health conditions affecting financial planning, relationship status, and intergenerational wealth intentions. When that information is processed by external AI model providers in plaintext, on their infrastructure, under their terms, the governance question is not hypothetical. It is immediate: who else, technically, can access what your clients told you in confidence? The answer should be: no one. The matrix tells you whether it is.

The critical regulatory context

SEC Investment Advisers Act & Fiduciary Duty AI-driven investment recommendations must be in the client's best interest, explainable, and auditable. The SEC is actively examining AI use in advisory services — suitability determinations, financial planning outputs, and client communications generated by AI carry the same fiduciary standard as human advice.

FINRA Rules (for broker-dealer affiliates) FINRA suitability and best interest standards (Regulation Best Interest) apply to AI-assisted investment recommendations. Supervisory system requirements extend to AI systems generating client recommendations.

SEC Regulation BI (Best Interest) For broker-dealers, Reg BI requires that AI-assisted recommendations be in the client's best interest with full disclosure of conflicts. AI systems generating product recommendations must satisfy the care, conflict, and disclosure obligations.

MiFID II (for EU clients) Suitability assessment, appropriateness testing, and record-keeping requirements apply to AI-assisted advisory for EU clients. AI-generated financial plans and investment recommendations must meet MiFID II documentation standards.

State Privacy Laws (CCPA, NY SHIELD, others) Client financial data processed by AI systems is subject to state privacy laws that may impose data residency, deletion, and disclosure obligations. Multi-state wealth managers face a patchwork of state-level requirements that AI governance frameworks must address.

Five questions for the board

- 1. Do we own our AI —** or are our clients' most sensitive financial details being processed on infrastructure where our model provider has ongoing access by design?
- 2. When our advisors use AI** to generate financial plans, suitability assessments, or client communications, do those outputs meet our fiduciary standard — and can we prove it with an audit trail?
- 3. If a client asked us** what AI systems processed their financial information, who had access to it, and what it was used for — could we answer completely and accurately?
- 4. Do our AI governance frameworks** satisfy SEC examination, FINRA supervision, Regulation BI, and MiFID II requirements simultaneously across our full advisory operation?
- 5. When our AI agents** conduct client onboarding, generate financial plans, or communicate with clients autonomously — are those interactions logged, auditable, and compliant with every applicable regulatory standard?

Completion: 60–90 min | Score: 0–100 | 25 control questions | information@institutional.ai

TABLE OF CONTENTS

SECTION	PAGE
For the Board — Executive Brief.....	2
Table of Contents	4
1. Executive Overview	5
2. The Wealth Management AI Governance Challenge	7
3. The Framework	12
<i>The Five AI Ecosystems</i>	12
<i>The Five Pillars of Control</i>	14
<i>How the Matrix Works</i>	16
4. The Assessment	18
<i>Ecosystem 1: Power</i>	18
<i>Ecosystem 2: Compute</i>	21
<i>Ecosystem 3: Data Centers</i>	24
<i>Ecosystem 4: Models</i>	27
<i>Ecosystem 5: Agents</i>	30
5. Scoring Summary & Heat Map	33
6. Score Interpretation	35
7. Sector Benchmarks	38
8. Critical Cell Analysis	43
9. Recommendations	47
10. Red Flags & Green Flags	57
11. What to Watch: The Next 12 Months	61
12. Final Guidance	64
13. About Institutional AI	67
Next Steps & Engagement.....	69
Glossary of Key Terms	71

1 EXECUTIVE OVERVIEW

What This Assessment Measures

The AI Sovereignty Assessment for Wealth Management measures your institution's verified ability to own, govern, and audit the AI systems that advise clients, assess suitability, generate financial plans, and manage the client relationship — across five governance control dimensions and five AI infrastructure layers.

The assessment produces a 5x5 matrix of 25 specific, answerable governance questions. A score of 1 to 4 per cell — maximum 100 total — reveals exactly which infrastructure-governance intersections are exposed. For wealth managers, every exposed cell represents a risk to something that is simultaneously a regulatory obligation, a fiduciary duty, and a client trust commitment: the sovereign protection of the most sensitive financial information your clients will ever share.

INSTITUTIONAL AI AI SOVEREIGNTY ASSESSMENT					
THE 5x5 CONTROL MATRIX · 25 GOVERNANCE INTERSECTIONS · MAX SCORE 100 · © 2026 INSTITUTIONAL AI					
FIVE PILLARS × five ecosystems	1 — POWER Energy infrastructure	2 — COMPUTE GPU / chip infrastructure	3 — DATA CENTERS Cloud & physical infrastructure	4 — MODELS LLMs & AI systems	5 — AGENTS Agentic applications
PILLAR 1 Jurisdictional Control Where does it execute, under which law?	Where is energy data processed and under which jurisdiction? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Under which jurisdiction does GPU compute actually execute? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Can you prove where every data center workload executes? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Which jurisdiction governs model training, storage, and serving — and does it grant access to weights? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Where do agent actions execute and where do decision logs reside? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 2 Logical Control Who can access it, when, with what proof?	Who has privileged access to energy systems — and is it logged in your SIEM? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Who accesses GPU clusters — including provider support engineers? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Can you prove no unauthorized access in the past 18 months? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Can your model provider access your queries and outputs — in your logs? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Who can approve or halt agents — are all actions logged in your systems? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 3 Technical Control Who holds the encryption keys?	Do you control encryption keys for energy and ESG data? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Who controls keys for AI training data and model weights? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do you hold HYOK for all sensitive data — or does your cloud provider? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Does your model provider process your data in plaintext on their infra? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do you have cryptographic controls over what agents can access and output? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 4 Operational Control Do you have real-time visibility?	Real-time visibility into energy use and carbon intensity per workload? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Real-time visibility into compute utilization and cost across all GPUs? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Can you detect a residency violation or breach within minutes, not hours? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Real-time visibility into model behavior, quality, and decision provenance? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Can you monitor, pause, or audit every agent action in real time? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 5 Contractual Control Do you have enforceable rights?	Do energy contracts include audit rights, portability, and exit provisions? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do compute agreements protect against unilateral capacity restrictions? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do data center agreements include audit rights and deletion certification? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Do you own your query logs and outputs — or does the provider retain? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Do agent agreements cover liability for autonomous decisions and data exit? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4

Critical cell — universally low governance; most institutions at Level 1 (Reactive) → 4 (Sovereign) per cell · Max total: 100 pts

The Client Privacy Imperative

Wealth management clients share information with their advisors that they share with no one else — not their accountants, not their attorneys, and frequently not their families. Net worth across all assets and liabilities. Tax position and planning strategies. Estate intentions and beneficiary designations. Family financial dynamics including disagreements, dependencies, and vulnerabilities. Health conditions affecting financial planning. Relationship status and its implications for wealth structures.

When that information is submitted to an external AI model — to generate a financial plan, to assess portfolio suitability, to draft a client communication, to synthesize financial advice — it is processed in plaintext on the provider's infrastructure under terms that most wealth management legal teams have not reviewed with the rigour those terms deserve. The question that every wealth management board should be asking is not whether they use AI. It is whether the AI they use processes client information with the same confidentiality standard their advisors apply.

The fiduciary standard does not end at the advisor relationship. It extends to every system that processes client information to inform advisory decisions. If that system is technically accessible to a third-party provider by design, the fiduciary obligation is not being met at the technical layer — regardless of what the contract says.

The Suitability and Explainability Obligation

AI systems contributing to investment recommendations, financial plans, and suitability assessments in wealth management carry regulatory obligations that are immediate and explicit. The SEC's fiduciary standard for RIAs, FINRA's suitability and Regulation Best Interest requirements for broker-dealers, and MiFID II's appropriateness and suitability standards for EU client interactions all require that AI-influenced advice be: in the client's best interest, explainable in terms the client can understand, and auditable to demonstrate that the process was sound.

The Operational × Models and Logical × Models cells directly determine whether you can satisfy these requirements. If your suitability AI operates without real-time monitoring, if the decision provenance of AI-generated financial plans is not captured in institution-controlled logs, and if your model provider has access to client financial profiles that are not tracked in your SIEM, your ability to demonstrate regulatory compliance in an examination is structurally compromised.

How to Score Each Cell

Level 1	Reactive No visibility or control. Relying entirely on provider assurances or standard contracts. Score: 1 point.
Level 2	Evolving Partial visibility or contractual controls only. Aware of the gap, some mitigating measures, but technical enforcement absent or incomplete. Score: 2 points.
Level 3	Governed Active monitoring, enforced contractual rights, partial technical controls. Demonstrable to regulators, but not cryptographically provable. Score: 3 points.
Level 4	Sovereign Full technical and contractual sovereignty. Cryptographically verifiable, continuously monitored, independently auditable. Provider assurances not relied upon. Score: 4 points.

2 THE WEALTH MANAGEMENT AI GOVERNANCE CHALLENGE

The AI Use Cases — and Their Governance Implications

Financial Planning and Goal-Based Advisory

AI-driven financial planning — retirement projections, education funding, cash flow analysis, goal-based portfolio construction — is being used at scale across RIAs, private banks, and wirehouses. The governance implication: financial plans generated by AI that inform client investment decisions carry the same fiduciary standard as human-generated plans. When a plan is based on incorrect assumptions, ignores relevant client information, or produces outputs that do not reflect the client's actual circumstances, the AI system's role in that failure must be reconstructable from audit logs. Most wealth management AI financial planning systems do not produce the decision provenance records required to satisfy that standard.

Primary matrix exposure: *Operational × Models. Logical × Models. Contractual × Models. Agents column for autonomous plan generation.*

Suitability Assessment and Portfolio Recommendation

AI systems are increasingly used to assess client suitability, match clients to investment products, and generate portfolio recommendations. The governance implication: Regulation BI, FINRA suitability rules, and MiFID II appropriateness requirements all apply to AI-assisted recommendations. The AI system's assessment of the client's risk tolerance, investment objectives, time horizon, and financial situation must be documentable, explainable, and demonstrably aligned with the client's best interest. Suitability AI that cannot produce these records creates direct regulatory and litigation exposure.

Primary matrix exposure: *Logical × Models (critical). Operational × Models. Contractual × Models. Operational × Agents for autonomous recommendation delivery.*

Client Communication and Relationship Management

AI-generated client communications — market commentary, portfolio updates, financial plan summaries, proactive outreach — are becoming standard across wealth management platforms. The governance implication: SEC Marketing Rule requirements apply to AI-generated client communications that reference performance or make investment claims. FINRA supervision requirements extend to AI systems generating client-facing content. An AI communication that misrepresents portfolio performance, makes unsuitable recommendations, or creates misleading impressions creates regulatory exposure that is indistinguishable from human-authored violations.

Primary matrix exposure: *Operational × Agents. Contractual × Agents. Logical × Models for content generation AI.*

Client Onboarding and KYC

AI-driven client onboarding — identity verification, KYC data collection, risk profiling, suitability questionnaire processing, account opening — is transforming the wealth management client acquisition process. The governance implication: AI systems processing client identity, financial profile, and risk assessment data create obligations under FinCEN AML requirements, FATF guidance, state KYC regulations, and GDPR for European clients. The auditability of AI-driven onboarding decisions — particularly adverse action decisions — is an explicit regulatory requirement.

Primary matrix exposure: *Jurisdictional × Models. Logical × Agents. Contractual × Agents. Technical × Data Centers.*

Estate Planning and Intergenerational Wealth

AI is being used to assist with estate planning analysis, beneficiary designation review, trust structure optimization, and intergenerational wealth transfer planning. The governance implication: estate planning AI processes information that is extraordinarily sensitive — family structures, relationship dynamics, health conditions, end-of-life intentions, and beneficiary conflicts. This information, if exposed through inadequate AI governance, creates privacy, fiduciary, and potentially legal consequences that extend beyond the primary client to their families and estates.

Primary matrix exposure: *Technical × Models (critical for data sensitivity). Jurisdictional × Data Centers. Contractual × Models.*

Tax Planning and Optimization

AI-driven tax loss harvesting, asset location optimization, tax-efficient withdrawal sequencing, and Roth conversion analysis are standard features of advanced wealth management platforms. The governance implication: tax planning AI processes tax return data, account structures, and future income projections that are among the most sensitive financial information a client possesses. This data submitted to external AI models creates privacy exposure under IRC confidentiality standards, state tax confidentiality laws, and general fiduciary duty.

Primary matrix exposure: *Technical × Models. Jurisdictional × Models. Logical × Models. Contractual × Models.*

Risk Profiling and Behavioral Finance

AI systems analyzing client behavioral patterns, risk tolerance indicators, and financial decision-making tendencies are being used to improve advisory quality and personalization. The governance implication: behavioral and psychographic data about clients is among the most sensitive information a wealth manager holds. AI systems processing behavioral data to profile clients must be governed under consumer protection, privacy, and fiduciary frameworks simultaneously. The use of behavioral AI outputs to influence product selection creates conflicts of interest that require explicit governance controls.

Primary matrix exposure: *Technical × Data Centers. Logical × Models. Contractual × Models. Jurisdictional × Models.*

Compliance Monitoring and Regulatory Reporting

AI-driven compliance surveillance — advisor supervision, regulatory reporting validation, complaint analysis, and conduct monitoring — is standard in large wealth management operations. The governance implication: compliance AI monitoring advisor conduct, client complaints, and regulatory reporting must itself be governed to the standard of the regulations it monitors. An AI compliance system that misses a violation due to model drift creates direct regulatory exposure.

Primary matrix exposure: *Operational × Models. Logical × Agents. Technical × Data Centers.*

The Regulatory Framework

Framework	Primary Applicability	Primary Matrix Impact
SEC Investment Advisers Act & Fiduciary Duty	All SEC-registered investment advisers. Fiduciary duty extends to AI-assisted advice — explainability, auditability, and client-best-interest standards apply to AI suitability and planning systems.	<i>Logical × Models. Operational × Models. Contractual × Models.</i>
Regulation Best Interest (Reg BI)	SEC-registered broker-dealers. Care, conflict, and disclosure obligations apply to AI-assisted investment recommendations. Supervisory system requirements extend to AI recommendation engines.	<i>Operational × Models. Logical × Agents. Contractual × Agents.</i>
FINRA Rules & Supervision Requirements	FINRA-member firms. Supervisory system requirements, suitability rules, and conduct standards apply to AI systems generating client recommendations and communications. FINRA examination focus on AI supervision is intensifying.	<i>Operational × Agents. Logical × Models. Contractual × Models.</i>
MiFID II (for EU clients)	Wealth managers serving EU clients. Suitability assessment, appropriateness testing, record-keeping, and best execution requirements apply to AI-assisted advisory. AI-generated financial plans must meet MiFID II documentation standards.	<i>Operational × Models. Contractual × Models. Jurisdictional × Data Centers.</i>
SEC Marketing Rule (Rule 206(4)-1)	All SEC-registered investment advisers. AI-generated client communications, performance references, and investment commentary must meet substantiation and accuracy requirements.	<i>Operational × Agents. Contractual × Models. Logical × Models.</i>
GDPR / UK GDPR (for European clients)	Wealth managers processing EU/UK resident client data. Data minimization,	<i>Jurisdictional × Models. Technical × Data Centers. Contractual × Models.</i>

	purpose limitation, data subject rights, and cross-border transfer restrictions apply to AI systems processing client financial profiles.	
State Privacy Laws (CCPA, NY SHIELD, others)	Multi-state wealth managers. A patchwork of state-level data privacy, breach notification, and financial data protection requirements apply to AI systems processing client financial information.	<i>Jurisdictional × Data Centers. Contractual × Models. Technical × Data Centers.</i>
FinCEN AML / BSA Requirements	All financial institutions including registered investment advisers. AI-driven KYC, client risk scoring, and suspicious activity monitoring must produce audit trails satisfying BSA record-keeping and SAR documentation requirements.	<i>Logical × Agents. Operational × Agents. Contractual × Models.</i>
ERISA (for retirement plan advisory)	Advisers to ERISA-governed retirement accounts. Fiduciary AI systems advising on retirement assets must satisfy ERISA's prudent expert standard — explainability and documentation requirements are explicit.	<i>Contractual × Models. Operational × Models. Logical × Agents.</i>
EU AI Act (developing)	AI systems used in investment advice and creditworthiness assessment may qualify as high-risk. Conformity assessment, human oversight, and transparency requirements for EU client-facing AI advice systems.	<i>Logical × Models. Operational × Models. Contractual × Models.</i>

The Trust Asymmetry Problem

Wealth management is built on an information asymmetry that clients accept in exchange for trust: the client shares everything, the advisor uses that information exclusively in the client's interest. AI governance failures do not break this trust abstractly — they break it specifically. A client whose estate planning information was processed by an AI provider that retained interaction logs did not consent to that retention. A client whose behavioral risk profile was used to optimize product recommendations through an AI system with inadequate conflict-of-interest controls was not served with the loyalty their fiduciary relationship entitles them to.

The matrix quantifies where that trust asymmetry is technically protected and where it is only contractually promised. The difference between Level 2 and Level 4 governance is the difference between a provider that promises not to access client information and a provider that technically cannot. For the clients who share their most sensitive financial details with your advisors, that distinction matters.

Wealth management clients do not read your AI governance policies. They experience the trust that your governance either earns or forfeits. The matrix determines which experience they actually receive.

3 THE FRAMEWORK

The Five AI Ecosystems in Wealth Management

ECOSYSTEM 1 — POWER *Energy infrastructure*

The energy infrastructure powering your wealth management platforms, client portals, advisor workstations, and data processing environments. For wealth managers, energy governance is primarily a sustainability and ESG reporting obligation — increasingly one that clients with sustainability mandates apply to their advisors' own operations.

Why it matters for wealth management: ESG-focused clients expect their wealth managers to demonstrate the same sustainability discipline they apply to their portfolios. Energy data governance is becoming a client reporting and due diligence requirement, not only a cost management function.

Representative providers: *Utility companies, renewable energy providers, data center energy platforms, cloud provider sustainability tools, carbon accounting and ESG reporting vendors.*

ECOSYSTEM 2 — COMPUTE *GPU / chip infrastructure*

The compute environments running your AI advisory systems, financial planning engines, suitability tools, client communication AI, and data processing infrastructure. Compute governance in wealth management is primarily about data protection during processing — ensuring client financial data processed by AI is not accessible to provider infrastructure teams.

Why it matters for wealth management: Client wealth data processed in cloud compute environments may be accessible to cloud provider support engineers under standard agreements. The sensitivity of wealth management client data — net worth, tax position, estate structure — makes logical control over compute access an acute fiduciary concern.

Representative providers: *AWS, Azure, Google Cloud, Oracle Cloud, IBM Cloud, specialized financial services cloud providers, on-premises compute for the most sensitive client data processing.*

ECOSYSTEM 3 — DATA CENTERS *Cloud & physical infrastructure*

The data centers and cloud storage where client financial profiles, household data, financial plans, suitability assessments, and advisor interaction records reside. For wealth managers, data center sovereignty determines the jurisdictional and technical protection of the most sensitive personal financial information your clients possess.

Why it matters for wealth management: Client financial data subject to GDPR for European clients, CCPA for California residents, and state privacy laws for multi-state operations creates data residency obligations that data center governance must technically enforce — not merely contractually assert. The data center layer is where most large wealth managers have the most mature governance, and where the foundation for Models and Agents governance must be built.

Representative providers: *AWS, Azure, Google Cloud, Equinix, Digital Realty, specialized financial services colocation, private cloud environments for the most sensitive client data.*

ECOSYSTEM 4 — MODELS *LLMs & AI systems*

The AI models processing client financial information and generating advisory outputs — financial planning AI, suitability assessment systems, portfolio recommendation engines, client communication AI, tax optimization models, and the foundation models underlying conversational advisor tools. This is the most acute governance gap in wealth management AI today.

Why it matters for wealth management: The data submitted to external AI models in wealth management is qualitatively different from almost any other industry. Estate planning scenarios, family relationship dynamics, health conditions, end-of-life financial intentions — this information is processed in plaintext by external model providers under standard API terms. The Technical × Models and Contractual × Models cells determine whether that processing is protected or exposed.

Representative providers: *Anthropic, OpenAI, Google, Microsoft (foundation model APIs for advisor AI tools), Salesforce Einstein, Microsoft Copilot for Financial Services, specialized wealth management AI platforms, proprietary planning and suitability AI systems.*

ECOSYSTEM 5 — AGENTS *Agentic applications*

The autonomous AI agents conducting client onboarding, generating financial plans, delivering portfolio updates, monitoring suitability on an ongoing basis, and communicating with clients through digital channels. Agentic AI in wealth management is moving from digital assistant to autonomous advisor — a transition that governance frameworks have not kept pace with.

Why it matters for wealth management: An agent that autonomously generates a financial plan, delivers it to a client, and recommends investment actions has performed an advisory function under applicable regulatory frameworks. The fiduciary, suitability, and disclosure obligations that apply to human advisors apply equally to the agent — and the audit trail requirements are identical. Most wealth management agent deployments cannot produce those audit trails.

Representative providers: *Salesforce Agentforce, Microsoft Copilot Studio, Google AgentSpace, proprietary robo-advisory and digital advisory frameworks, wealth management fintech agent vendors, AI-powered client portal platforms.*

The Five Pillars of Control

PILLAR 1 — JURISDICTIONAL CONTROL *Where does it execute, under which law?*

Proving the physical and legal location of every AI workload processing client financial data — with technical evidence. For wealth managers serving clients in multiple jurisdictions, jurisdictional control determines which laws govern client data and who can compel access to it.

Wealth management implication: GDPR creates explicit data residency obligations for EU client data. CCPA and state privacy laws create residency and disclosure obligations for US clients. Multi-jurisdictional wealth managers face a patchwork of requirements that only technical jurisdictional enforcement — not contractual commitments — can reliably satisfy.

PILLAR 2 — LOGICAL CONTROL *Who can access it, when, with what proof?*

Proving who accessed client financial information, when, and under what conditions — with immutable, independently verifiable evidence. In wealth management, logical control extends to advisor access, platform access, model provider infrastructure access, and agent-initiated data access.

Wealth management implication: SEC examination staff increasingly request access logs for systems processing client data. FINRA supervisory requirements mandate records of who accessed client information and when. When a data breach occurs or a fiduciary violation is alleged, the logical access record is the first evidence regulators demand. Most wealth management AI systems do not produce that record in institution-controlled systems.

PILLAR 3 — TECHNICAL CONTROL *Who holds the encryption keys?*

Cryptographic protection of client financial data across all five ecosystems. BYOK gives you key management but the provider can still access data during AI model processing. HYOK combined with confidential computing means the provider technically cannot access client estate plans, tax positions, or family financial dynamics during model inference.

Wealth management implication: For wealth management, the sensitivity of client data makes the distinction between contractual and technical protection especially significant. Client estate planning information, health-related financial planning, and family wealth dynamics are categories of information that clients share with their advisors specifically because of the expectation of exclusive confidentiality. Technical controls enforce that expectation; contractual promises assert it.

PILLAR 4 — OPERATIONAL CONTROL *Do you have real-time visibility?*

Real-time visibility into AI model behavior, suitability output quality, and advisory decision provenance across all AI systems contributing to client recommendations and financial plans.

Wealth management implication: A suitability AI that drifts from its validated behavior without triggering operational monitoring can produce recommendations that no longer reflect the client's actual financial situation. A financial planning AI that degrades in quality without detection can deliver plans that no longer meet the prudent expert standard. Real-time operational control is not just compliance infrastructure — it is the mechanism by which you ensure advice quality in a world where AI systems can change without human awareness.

PILLAR 5 — CONTRACTUAL CONTROL *Do you have enforceable rights?*

Enforceable legal rights to audit AI providers, exit agreements on your terms, and protect client financial information in the AI relationship. Standard model API contracts do not provide the rights required for wealth management fiduciary, regulatory, and client protection obligations.

Wealth management implication: Your model API agreements almost certainly do not provide: explicit ownership of client interaction logs; certified deletion of client financial data submitted as queries; audit rights exercisable without provider consent; or exit provisions ensuring complete client data deletion within timelines consistent with your data governance obligations. Each gap creates direct fiduciary and regulatory exposure.

How the Matrix Works

INSTITUTIONAL AI AI SOVEREIGNTY ASSESSMENT THE 5x5 CONTROL MATRIX · 25 GOVERNANCE INTERSECTIONS · MAX SCORE 100 · © 2026 INSTITUTIONAL AI					
FIVE PILLARS × five ecosystems	1 — POWER Energy infrastructure	2 — COMPUTE GPU / chip infrastructure	3 — DATA CENTERS Cloud & physical infrastructure	4 — MODELS LLMs & AI systems	5 — AGENTS Agentic applications
PILLAR 1 Jurisdictional Control Where does it execute, under which law?	Where is energy data processed and under which jurisdiction? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Under which jurisdiction does GPU compute actually execute? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Can you prove where every data center workload executes? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Which jurisdiction governs model training, storage, and serving — and does it grant access to weights? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Where do agent actions execute and where do decision logs reside? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 2 Logical Control Who can access it, when, with what proof?	Who has privileged access to energy systems — and is it logged in your SIEM? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Who accesses GPU clusters — including provider support engineers? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Can you prove no unauthorized access in the past 18 months? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Can your model provider access your queries and outputs — in your logs? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Who can approve or halt agents — are all actions logged in your systems? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 3 Technical Control Who holds the encryption keys?	Do you control encryption keys for energy and ESG data? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Who controls keys for AI training data and model weights? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do you hold HYOK for all sensitive data — or does your cloud provider? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Does your model provider process your data in plaintext on their infra? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do you have cryptographic controls over what agents can access and output? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 4 Operational Control Do you have real-time visibility?	Real-time visibility into energy use and carbon intensity per workload? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Real-time visibility into compute utilization and cost across all GPUs? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Can you detect a residency violation or breach within minutes, not hours? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Real-time visibility into model behavior, quality, and decision provenance? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Can you monitor, pause, or audit every agent action in real time? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4
PILLAR 5 Contractual Control Do you have enforceable rights?	Do energy contracts include audit rights, portability, and exit provisions? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do compute agreements protect against unilateral capacity restrictions? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	Do data center agreements include audit rights and deletion certification? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Do you own your query logs and outputs — or does the provider retain? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4	CRITICAL CELL Do agent agreements cover liability for autonomous decisions and data exit? <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4

Critical cell — universally low governance; most institutions at Level 1 (Reactive) → 4 (Sovereign) per cell · Max total: 100 pts

Scoring principle: Score each cell based on demonstrable evidence — not what you believe, what providers have attested, or what contracts assert. For wealth managers, the most important honesty is in the Models column: this is where client financial data is most acutely exposed, and it is where most institutions will find Level 1 scores that require immediate attention.

4 THE ASSESSMENT

Complete all 25 questions. Score each cell 1–4 based on demonstrable evidence. Recommended team: CTO, CISO, General Counsel, Chief Compliance Officer, Head of Advisory Technology, Head of Client Services. Time: 60–90 minutes.

Pay particular attention to the Models column — this is where client financial data is most acutely exposed in most wealth management operations. A Level 1 score in Logical × Models means your model provider can access client estate plans, tax positions, and family wealth dynamics that your clients shared in confidence.

ECOSYSTEM 1 — POWER | Energy infrastructure

PILLAR 1 — JURISDICTIONAL CONTROL | Where does it execute, under which law?

Where is the energy data powering your wealth management platforms processed — and under which jurisdiction, satisfying your obligations to clients who may reside in multiple regulatory regimes simultaneously?

- Level 1** Reactive — no visibility; relying on provider assurances
- Level 2** Evolving — partial contractual controls; limited technical enforcement
- Level 3** Governed — active monitoring, contractual rights, partial technical controls
- Level 4** Sovereign — cryptographically verifiable, continuously monitored, independently auditable

P1×E1 Score: _____ / 4

Notes: _____

PILLAR 2 — LOGICAL CONTROL | Who can access it, when, with what proof?

Who has privileged access to the infrastructure systems powering your wealth management platforms — and is that access immutably logged to your own SIEM in a format producible for regulatory examination?

- Level 1** Reactive — no visibility; relying on provider assurances
- Level 2** Evolving — partial contractual controls; limited technical enforcement
- Level 3** Governed — active monitoring, contractual rights, partial technical controls
- Level 4** Sovereign — cryptographically verifiable, continuously monitored, independently auditable

P2×E1 Score: _____ / 4

Notes: _____

PILLAR 3 — TECHNICAL CONTROL | Who holds the encryption keys?

Do you control the encryption keys for your energy management and infrastructure telemetry data?

- Level 1** Reactive — no visibility; relying on provider assurances
- Level 2** Evolving — partial contractual controls; limited technical enforcement
- Level 3** Governed — active monitoring, contractual rights, partial technical controls
- Level 4** Sovereign — cryptographically verifiable, continuously monitored, independently auditable

P3×E1 Score: _____ / 4

Notes: _____

PILLAR 4 — OPERATIONAL CONTROL | Do you have real-time visibility?

Do you have real-time visibility into energy consumption per platform — sufficient to satisfy ESG reporting obligations and client sustainability expectations that increasingly include infrastructure accountability?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P4×E1 Score: _____ / 4

Notes: _____

PILLAR 5 — CONTRACTUAL CONTROL | *Do you have enforceable rights?*

Do your energy and infrastructure contracts include audit rights and exit provisions that satisfy your operational resilience obligations under SEC, FINRA, FCA, and applicable state regulatory frameworks?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P5×E1 Score: _____ / 4

Notes: _____

ECOSYSTEM 1 — POWER SUBTOTAL: _____ / 20

ECOSYSTEM 2 — COMPUTE | GPU / chip infrastructure

PILLAR 1 — JURISDICTIONAL CONTROL | *Where does it execute, under which law?*

Under which jurisdiction do the compute resources running your AI-driven financial planning, portfolio advisory, and client suitability systems actually execute?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P1×E2 Score: _____ / 4

Notes: _____

PILLAR 2 — LOGICAL CONTROL | *Who can access it, when, with what proof?*

Who has access to the compute environments processing client wealth data and AI advisory outputs — including cloud provider support engineers — and is that access logged in your systems?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P2×E2 Score: _____ / 4

Notes: _____

PILLAR 3 — TECHNICAL CONTROL | *Who holds the encryption keys?*

Who controls the encryption keys for the AI model weights, training data, and client behavioral datasets underlying your suitability and financial planning AI systems?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P3×E2 Score: _____ / 4

Notes: _____

PILLAR 4 — OPERATIONAL CONTROL | *Do you have real-time visibility?*

Do you have real-time visibility into compute utilization, cost, and performance across all AI infrastructure supporting your advisory, planning, and client engagement systems?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P4×E2 Score: _____ / 4

Notes: _____

PILLAR 5 — CONTRACTUAL CONTROL | *Do you have enforceable rights?*

Do your compute and cloud agreements protect client data sovereignty and include workload portability rights sufficient for migration without client data exposure during transition?

- Level 1** *Reactive — no visibility; relying on provider assurances*

- Level 2 *Evolving — partial contractual controls; limited technical enforcement*
- Level 3 *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4 *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P5×E2 Score: _____ / 4

Notes: _____

ECOSYSTEM 2 — COMPUTE SUBTOTAL: _____ / 20

ECOSYSTEM 3 — DATA CENTERS | *Cloud & physical infrastructure*

PILLAR 1 — JURISDICTIONAL CONTROL | *Where does it execute, under which law?*

Can you demonstrate with technical evidence where every client financial profile, household wealth data, and AI-generated advice record executes — satisfying SEC, FCA, MiFID II, and applicable state regulatory data residency obligations?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P1×E3 Score: _____ / 4

Notes: _____

PILLAR 2 — LOGICAL CONTROL | *Who can access it, when, with what proof?*

Can you prove that no unauthorized person accessed any client financial profile, wealth plan, or AI-generated suitability assessment in the past 18 months — satisfying SEC, FINRA, FCA, and fiduciary audit requirements?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P2×E3 Score: _____ / 4

Notes: _____

PILLAR 3 — TECHNICAL CONTROL | *Who holds the encryption keys?*

Do you hold your own encryption keys (HYOK) for all client financial data — net worth, tax records, estate documents, family financial structures — or does your cloud provider manage and control them?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P3×E3 Score: _____ / 4

Notes: _____

PILLAR 4 — OPERATIONAL CONTROL | *Do you have real-time visibility?*

Can you detect an unauthorized access to client financial data, a suitability AI anomaly, or a data residency violation within minutes — satisfying your operational resilience obligations and client data protection commitments?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P4×E3 Score: _____ / 4

Notes: _____

PILLAR 5 — CONTRACTUAL CONTROL | *Do you have enforceable rights?*

Do your data center agreements include unlimited audit rights, subprocessor transparency with client notification rights, certified deletion of client financial data, and exit terms exercisable within timelines consistent with client data protection obligations?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P5×E3 Score: _____ / 4

Notes: _____

ECOSYSTEM 3 — DATA CENTERS SUBTOTAL: _____ / 20

ECOSYSTEM 4 — MODELS | LLMs & AI systems

PILLAR 1 — JURISDICTIONAL CONTROL | *Where does it execute, under which law?*

Under which jurisdiction are the AI models processing your clients' most sensitive financial data — net worth, tax position, estate structure, family relationships — trained, stored, and served?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P1×E4 Score: _____ / 4

Notes: _____

PILLAR 2 — LOGICAL CONTROL | *Who can access it, when, with what proof?*

Can your AI model providers access the client financial data your models process — including net worth, tax information, estate structures, and family financial dynamics — and is that access logged in your own infrastructure?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P2×E4 Score: _____ / 4

Notes: _____

PILLAR 3 — TECHNICAL CONTROL | *Who holds the encryption keys?*

When your AI models process client financial profiles, suitability assessments, or wealth planning scenarios, does the model provider process this extraordinarily sensitive data in plaintext on their own infrastructure?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P3×E4 Score: _____ / 4

Notes: _____

PILLAR 4 — OPERATIONAL CONTROL | *Do you have real-time visibility?*

Do you have real-time visibility into AI model behavior, suitability output quality, and advisory decision provenance across all AI systems contributing to client recommendations and financial plans?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P4×E4 Score: _____ / 4

Notes: _____

PILLAR 5 — CONTRACTUAL CONTROL | *Do you have enforceable rights?*

Do your AI model provider agreements give you explicit rights over the client financial data processed, interaction logs, suitability outputs, and wealth planning scenarios generated — or does the provider retain residual rights over your clients' most sensitive financial information?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P5×E4 Score: _____ / 4

Notes: _____

ECOSYSTEM 4 — MODELS SUBTOTAL: _____ / 20

ECOSYSTEM 5 — AGENTS | *Agentic applications*

PILLAR 1 — JURISDICTIONAL CONTROL | *Where does it execute, under which law?*

Under which jurisdiction do your autonomous client onboarding, financial planning, and advisory communication agents execute — and where do their decision logs and client interaction records reside?

- Level 1 *Reactive — no visibility; relying on provider assurances*
- Level 2 *Evolving — partial contractual controls; limited technical enforcement*
- Level 3 *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4 *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P1×E5 Score: _____ / 4

Notes: _____

PILLAR 2 — LOGICAL CONTROL | *Who can access it, when, with what proof?*

Who can approve, modify, or halt your autonomous client advisory, onboarding, and financial planning agents — and is every agent interaction with client data immutably logged in institution-controlled systems?

- Level 1 *Reactive — no visibility; relying on provider assurances*
- Level 2 *Evolving — partial contractual controls; limited technical enforcement*
- Level 3 *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4 *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P2×E5 Score: _____ / 4

Notes: _____

PILLAR 3 — TECHNICAL CONTROL | *Who holds the encryption keys?*

Do you have cryptographic enforcement over what your autonomous advisory and onboarding agents can access, record, and transmit — protecting client financial privacy beyond policy rules that agents could circumvent?

- Level 1 *Reactive — no visibility; relying on provider assurances*
- Level 2 *Evolving — partial contractual controls; limited technical enforcement*
- Level 3 *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4 *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P3×E5 Score: _____ / 4

Notes: _____

PILLAR 4 — OPERATIONAL CONTROL | *Do you have real-time visibility?*

Can you monitor, pause, or audit every autonomous agent conducting client onboarding, financial planning, or advisory communications — in real time, with complete interaction logs accessible to regulators and clients?

- Level 1 *Reactive — no visibility; relying on provider assurances*
- Level 2 *Evolving — partial contractual controls; limited technical enforcement*
- Level 3 *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4 *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P4×E5 Score: _____ / 4

Notes: _____

PILLAR 5 — CONTRACTUAL CONTROL | *Do you have enforceable rights?*

Do your agentic AI vendor agreements include institutional liability for autonomous client advisory decisions, audit rights over every agent-client interaction exercisable by both you and regulators, and data exit provisions covering all client financial data processed by agents?

- Level 1** *Reactive — no visibility; relying on provider assurances*
- Level 2** *Evolving — partial contractual controls; limited technical enforcement*
- Level 3** *Governed — active monitoring, contractual rights, partial technical controls*
- Level 4** *Sovereign — cryptographically verifiable, continuously monitored, independently auditable*

P5×E5 Score: _____ / 4

Notes: _____

ECOSYSTEM 5 — AGENTS SUBTOTAL: _____ / 20

5 SCORING SUMMARY & HEAT MAP

Transfer your 25 cell scores below. Total each row and column. For wealth managers, the Models column score is the most important single indicator — it reveals the degree of technical protection your clients' most sensitive financial information receives during AI processing.

PILLAR	POWER	COMPUTE	DATA CENTERS	MODELS	AGENTS	TOTAL
JURISDICTIONAL CONTROL	___	___	___	___	___	___ / 20
LOGICAL CONTROL	___	___	___	___	___	___ / 20
TECHNICAL CONTROL	___	___	___	___	___	___ / 20
OPERATIONAL CONTROL	___	___	___	___	___	___ / 20
CONTRACTUAL CONTROL	___	___	___	___	___	___ / 20
ECOSYSTEM TOTAL	___ / 20	___ / 20	___ / 20	___ / 20	___ / 20	___ / 100

Reading Your Results

- Models column (max 20)** The most important column. If this scores below 10, your clients' estate planning information, tax positions, and family wealth dynamics are being processed by external AI without adequate technical controls or contractual protections.
- Agents column (max 20)** If this scores below 8 and you are deploying agentic AI for client advisory, onboarding, or financial planning, your autonomous client interactions lack the audit trails required by SEC, FINRA, MiFID II, and ERISA.
- Contractual row (max 20)** This determines whether the fiduciary protections you believe you have are actually enforceable. A score below 10 in the Contractual row means client data protections are aspirational rather than legally defensible.
- Data Centers vs Models gap** Compare these two column totals. Every point the Models column falls below the Data Centers column represents client financial data that is better protected at rest than when your AI is processing it.
- The fiduciary threshold** A total score below 40 represents a governance posture that is difficult to defend in SEC or FINRA examination. A score below 25 represents a posture that client due diligence processes will increasingly identify as inadequate.

Notes: _____

TOTAL SOVEREIGNTY SCORE: _____ / 100