

THE AI SOVEREIGNTY ASSESSMENT

RETIREMENT PLAN PROVIDERS &
THIRD-PARTY ADMINISTRATORS

Should Your Institution Build, Rent, or Compose AI Infrastructure?

*Governing AI Sovereignty Across Plan Administration, Participant
Recordkeeping, Compliance, and Retirement Income*

A Strategic Decision Framework for Retirement Services Leaders

FOR THE BOARD | EXECUTIVE BRIEF |
RETIREMENT PLAN PROVIDERS & TPA | CONFIDENTIAL

How This Assessment Works — Two Instruments, One Program

This document contains two complementary instruments that work together. Understanding the relationship between them is essential before your leadership team begins.

INSTRUMENT 1

The 0–160 Strategic Assessment

Sections 1–4 of this document

Answers: What should our AI infrastructure strategy be?

This is a forward-looking strategic decision tool. It evaluates your regulatory environment, how important AI is to your ERISA fiduciary mission, your risk tolerance, and your financial capacity. The output maps to one of four strategies:

- **RENT** (score 0–40)
- **RENT + GOVERN** (score 41–80)
- **COMPOSE** (score 81–120)
- **BUILD** (score 121–160)

This tells you where you need to go.

INSTRUMENT 2

The 5×5 Control Matrix

Section 6 of this document

Answers: Where is our governance today?

This is a current-state diagnostic. It scores 25 specific governance intersections — five control pillars applied across five AI infrastructure ecosystems — based on what technical and contractual controls you actually have in place right now.

Each of the 25 cells is scored 1 (Reactive) to 4 (Sovereign). Maximum total: 100 points.

This tells you where you are standing today.

THE RELATIONSHIP BETWEEN THE TWO SCORES

Together they define your governance program.

The 0–160 score gives you your destination. The 5×5 matrix gives you your starting point. The distance between them is the work.

A CONCRETE EXAMPLE

A large TPA completes the 0–160 assessment and scores 95. That score falls in the 81–120 band — it says Compose is the right strategy. The institution needs hybrid sovereign architecture with a protected core for participant data.

They then complete the 5×5 Control Matrix and score 38 out of 100. That score says their current governance posture is essentially Reactive — most of the 25 cells are at Level 1.

The gap between 'you need to Compose' and 'you currently govern at Reactive level' is the entire work program. And the matrix tells you specifically which of the 25 cells are at Level 1 — those are your priorities. In retirement services, the Models and Agents columns are almost always at Level 1, meaning participant Social Security numbers and benefit data are being processed without the technical controls ERISA's prudent expert standard requires.

THE FOUR ASSESSMENT DIMENSIONS

Section 1	Regulatory & Fiduciary Requirements <i>How binding are your ERISA and DOL compliance obligations?</i>	40 pts
Section 2	Strategic Importance of AI <i>How existential is AI to your plan administration mission?</i>	40 pts
Section 3	Risk Tolerance & Sovereignty <i>What level of third-party dependency is acceptable under ERISA?</i>	40 pts
Section 4	Financial & Operational Capacity <i>Can you fund and operate sovereign AI infrastructure?</i>	40 pts

WHY THIS MATTERS NOW FOR RETIREMENT SERVICES

92% of advanced AI chips from a single company in a geopolitically contested region	70% of global AI compute capacity controlled by just five providers	78% of enterprises run mission-critical AI on infrastructure they cannot audit	2030 AI operations will consume 945 TWh annually — more electricity than Germany
---	---	--	--

The question is no longer whether to adopt AI — it is whether you can adopt AI without surrendering your ERISA fiduciary obligations to the infrastructure providers that run it.

FOUR OUTCOME STRATEGIES

0–40	RENT	Managed cloud with standard governance. AI is supplementary; commercial controls sufficient for current obligations.
41–80	RENT + GOVERN	Enhanced contractual controls and BYOK encryption. AI is important; contractual oversight with ERISA-specific terms adequate.
81–120	COMPOSE	Hybrid sovereign architecture. AI is strategic; own the critical participant data layer, rent the rest with strict governance.
121–160	BUILD	Full sovereign infrastructure. AI is existential; complete institutional independence required by ERISA and DOL obligations.

FIVE QUESTIONS FOR THE BOARD

1. Can management prove — with technical evidence — where every AI workload processing participant Social Security numbers and retirement savings data executes today?
2. If our primary AI provider restricted access for 90 days, what would happen to plan administration, compliance testing, and participant services?
3. Could we produce a complete AI decision audit trail for any participant benefit determination from 18 months ago within 24 hours — for DOL examination?
4. Do our AI governance frameworks meet the standard ERISA's prudent expert rule and DOL examination requirements demand?
5. Do our model provider agreements give us — and plan sponsors — the rights over participant data that ERISA Section 408(b)(2) requires?

TABLE OF CONTENTS

SECTION	PAGE
For the Board — Executive Brief.....	2
Table of Contents	4
1. Executive Overview	6
2. Understanding the Five Pillars of Control.....	9
<i>Pillar 1 — Jurisdictional Control.....</i>	10
<i>Pillar 2 — Logical Control</i>	10
<i>Pillar 3 — Technical Control</i>	11
<i>Pillar 4 — Operational Control.....</i>	11
<i>Pillar 5 — Contractual Control.....</i>	12
3. The Retirement Services AI Governance Challenge	13
4. The Assessment	17
<i>Section 1: Regulatory & Fiduciary Requirements.....</i>	18
<i>Section 2: Strategic Importance of AI.....</i>	20
<i>Section 3: Risk Tolerance & Sovereignty</i>	22
<i>Section 4: Financial & Operational Capacity</i>	24
5. Total Assessment Score & Interpretation	26
6. The 5×5 Control Matrix	28
7. Sector Benchmarks	40
8. Critical Cell Analysis	44
9. Detailed Recommendations & Strategic Playbook.....	47
<i>Score 0–40: Rent.....</i>	48
<i>Score 41–80: Rent + Govern.....</i>	49
<i>Score 81–120: Compose.....</i>	50
<i>Score 121–160: Build</i>	51
10. Migration Paths & Timelines	62
11. Red Flags & Green Flags.....	64
12. What to Watch: The Next 12 Months	67
13. Final Guidance.....	69
14. About Institutional AI.....	81
Next Steps & Engagement.....	83
Glossary of Key Terms	86
Legal Disclaimer	87

1 EXECUTIVE OVERVIEW

Purpose of This Assessment

The AI Sovereignty Assessment for Retirement Plan Providers and Third Party Administrators is a comprehensive diagnostic tool designed to help retirement services leadership determine the appropriate AI infrastructure strategy and understand exactly where their current governance stands. It contains two complementary instruments that answer two different but inseparable questions.

INSTRUMENT 1 — The 0–160 Strategic Assessment (Sections 1–4)

The 0–160 assessment answers: what should our AI infrastructure strategy be? It is a forward-looking strategic decision tool. It evaluates your regulatory environment and ERISA fiduciary obligations, how critical AI is to your plan administration mission, your risk tolerance for third-party dependency, and your financial and operational capacity. The output — Rent, Rent + Govern, Compose, or Build — tells you where you need to go.

0–40	RENT	Managed cloud with ERISA governance overlay. AI is supplementary.
41–80	RENT + GOVERN	Enhanced ERISA contractual controls and BYOK encryption.
81–120	COMPOSE	Hybrid sovereign architecture. Own the participant data layer, rent the rest.
121–160	BUILD	Full sovereign infrastructure. Complete independence for ERISA fiduciary AI.

INSTRUMENT 2 — The 5×5 Control Matrix (Section 6)

The 5×5 Control Matrix answers: where is our governance today? It is a current-state diagnostic. It scores 25 specific governance intersections — five pillars of control applied independently to each of five AI infrastructure ecosystems — based on what technical and contractual controls you actually have in place right now. Each cell is scored 1 (Reactive) to 4 (Sovereign). Maximum total: 100 points. This tells you where you are standing.

INSTITUTIONAL AI AI SOVEREIGNTY ASSESSMENT THE 5×5 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 least score 1 (Reactive) → 4 (Sovereign) per cell · Max total: 100 pts

TOGETHER THEY DEFINE YOUR GOVERNANCE PROGRAMME
The 0–160 score gives you your destination. The 5×5 matrix gives you your starting point. The distance between them is the work program.

WORKED EXAMPLE

A large TPA completes the 0–160 assessment and scores 95. That places them in the 81–120 band — Compose is the right strategy. They need a hybrid sovereign architecture with a protected core for participant data, particularly the Models ecosystem where participant SSNs are processed.

They then complete the 5×5 Control Matrix and score 38 out of 100. That says their current governance posture is essentially Reactive — most of the 25 cells sit at Level 1.

The gap between 'the institution needs to Compose' and 'it currently governs at Reactive level' is the entire work program. The matrix makes that gap specific: in retirement services, the Models and Agents columns are almost universally at Level 1. This means participant Social Security numbers, beneficiary designations, and hardship documentation are being processed without the technical controls ERISA's prudent expert standard actually requires. Those Level 1 cells in the Models and Agents columns are the program — not theoretical risk, but specific governance gaps with specific remediation steps and ERISA consequences.

Why This Assessment Matters for Retirement Services

ERISA	SSNs	DOL
The highest fiduciary standard in US law — applies fully to AI systems performing plan administration functions	Participant Social Security numbers — master key to identity theft — processed by AI under standard API terms at most providers	Examination focus on AI in retirement plan administration is intensifying — EBSA has broad authority under ERISA Section 504

The landscape for retirement services AI has fundamentally changed. AI is no longer supplementary to plan administration — it is embedded in compliance testing, retirement income projections, participant engagement, managed account services, enrollment automation, and benefit determination. Each of these functions carries ERISA's fiduciary standard to the AI layer that performs them.

The question is no longer whether to adopt AI in retirement services — it is whether you can adopt AI without surrendering your ERISA fiduciary obligations to the infrastructure providers that operate it.

What You Will Learn

- Quantify your AI sovereignty posture with a precise total score (0–160) and matrix score (0–100)
- Identify critical governance gaps across five control pillars and five AI infrastructure ecosystems
- Benchmark against peer retirement services institutions — large recordkeepers, regional TPAs, insurance platforms, government plan administrators
- Receive tailored recommendations for Build / Rent + Govern / Compose strategies aligned to your ERISA obligations
- Develop an actionable roadmap with specific next steps, timelines, and investment guidance
- Understand the ERISA fiduciary implications of your current AI governance posture for DOL examination readiness

How to Use This Assessment

Step 1	Complete Sections 1–4 honestly — score what you can demonstrate with evidence, not what you believe or what providers have attested
Step 2	Calculate your total score (maximum 160 points) and review your score-based strategy recommendation
Step 3	Complete the 5×5 Control Matrix (Section 6) to identify your specific governance exposure by infrastructure layer
Step 4	Review the Critical Cell Analysis (Section 8) for retirement-specific ERISA implications of your lowest-scoring cells
Step 5	Develop your action plan using the Detailed Recommendations and Migration Path sections
Step 6	Schedule a confidential briefing with Institutional AI for peer benchmarking and ERISA governance program design

Estimated completion time: 45–60 minutes