Institutional AI

# THE AI SOVEREIGNTY ASSESSMENT

## PRIVATE EQUITY (PARTIAL)EDITION

*Should Your Firm Build, Rent, or Compose AI Infrastructure?*

*A Strategic Decision Framework for Private Equity Leadership*

SEC ADVISERS ACT · LP FIDUCIARY · ILPA STANDARDS · M&A DUE DILIGENCE · DEAL INTELLIGENCE

APRIL 2026 | BOSTON, MA | GLOBAL ADVISORY SERVICES | www.institutionalai.net

## How This Assessment Works — Two Instruments, One Program

This document contains two complementary instruments that work together. The 0–160 Strategic Assessment tells you where you need to go. The 5×5 Control Matrix tells you where your AI governance actually stands today. Together they define the program.

INSTRUMENT 1

### The 0–160 Strategic Assessment

*Sections 1–4 of this document*

**Answers: What should our AI infrastructure strategy be?**

Evaluates your regulatory environment, deal intelligence sensitivity, risk tolerance, and financial capacity. 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?**

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 scored 1 (Reactive) to 4 (Sovereign). Maximum total: 100 points.

*This tells you where you are standing today.*

**THE GAP BETWEEN THE TWO SCORES IS YOUR PROGRAMME**

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

**A CONCRETE EXAMPLE**

A mid-market PE firm completes the 0–160 strategic assessment and scores 88 — Compose is the right strategy. They need hybrid sovereign architecture with a protected core for deal intelligence and LP data.

They complete the 5×5 Control Matrix and score 34 out of 100 — their current posture is Reactive. Most of the 25 cells are at Level 1, particularly in the Models and Agents columns where proprietary deal intelligence is being processed on external AI infrastructure under standard API terms.

**The gap between 'we need to Compose' and 'we currently govern at Reactive level' is the entire work program. The matrix identifies the specific Level 1 cells — the Models column first, then Agents — as priorities. The benchmark confirms urgency: they are below the peer range for firms their size. The strategic assessment sets the destination.**

**FIVE QUESTIONS FOR THE PARTNERSHIP**

1. Does every AI system processing proprietary deal intelligence, portfolio company data, and LP information execute on infrastructure the firm controls — and can management prove it with technical evidence?
2. When the firm submits deal pipeline data, investment thesis logic, or portfolio company financials to an external AI model, does the provider have ongoing technical access to that information by design?
3. If our primary AI provider restricted access for 90 days, what would happen to deal sourcing, due diligence, portfolio monitoring, and LP reporting?
4. Do our AI model provider agreements give the firm — and by extension our LPs — explicit rights over deal intelligence processed, diligence logs, and investment outputs?
5. When our AI agents process deal screening, portfolio monitoring, and LP reporting — are those actions completely logged in firm-controlled systems accessible to our GPs and regulators?

 

# TABLE OF CONTENTS (PARTIAL)



 

# 1 EXECUTIVE OVERVIEW

## The Deal Intelligence Problem

Your competitive advantage lives in what you see before others do. The proprietary deal signals, the pattern recognition built over years, the sourcing relationships, the operational playbooks, the investment thesis logic — the accumulated intelligence that differentiates your firm from every other GP competing for the same assets.

When that intelligence is submitted to an external AI model for deal screening, due diligence synthesis, portfolio company analysis, or LP reporting — it is processed on the provider's infrastructure under standard API terms that do not protect competitive deal intelligence. The provider has ongoing technical access to your pipeline, your thesis, your analysis, and your portfolio data. Every query is a window into how your firm thinks.

> *This is not a hypothetical risk. It is the operational reality of how AI models are served — and it is happening at most PE firms today, in every deal screening query and every portfolio monitoring report that touches an external AI model.*

## The Two-Instrument Framework

This assessment contains two complementary instruments that answer two different but inseparable questions.

### INSTRUMENT 1 — The 0–160 Strategic Assessment

Answers: what should our AI infrastructure strategy be? Evaluates regulatory obligations under the SEC Advisers Act, how critical AI is to deal sourcing and portfolio management, risk tolerance for third-party dependency on deal intelligence, and financial and operational capacity. Output: Rent, Rent + Govern, Compose, or Build.

### INSTRUMENT 2 — The 5×5 Control Matrix

Answers: where is our governance today? Scores 25 specific governance intersections across five control pillars and five AI infrastructure ecosystems. Each cell scored 1 (Reactive) to 4 (Sovereign). Maximum: 100 points.

 

**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? ☐1 ☐2 ☐3 ☐4 | Under which jurisdiction does GPU compute actually execute? ☐1 ☐2 ☐3 ☐4 | Can you prove where every data center workload executes? ☐1 ☐2 ☐3 ☐4 | Which jurisdiction governs model training, storage, and serving — and does it grant access to weights? ☐1 ☐2 ☐3 ☐4 | Where do agent actions execute and where do decision logs reside? ☐1 ☐2 ☐3 ☐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? ☐1 ☐2 ☐3 ☐4 | Who accesses GPU clusters — including provider support engineers? ☐1 ☐2 ☐3 ☐4 | Can you prove no unauthorized access in the past 18 months? ☐1 ☐2 ☐3 ☐4 | ☐ CRITICAL CELL Can your model provider access your queries and outputs — in your logs? ☐1 ☐2 ☐3 ☐4 | Who can approve or halt agents — are all actions logged in your systems? ☐1 ☐2 ☐3 ☐4 |
| **PILLAR 3 Technical Control** Who holds the encryption keys? | Do you control encryption keys for energy and ESG data? ☐1 ☐2 ☐3 ☐4 | Who controls keys for AI training data and model weights? ☐1 ☐2 ☐3 ☐4 | Do you hold HYOK for all sensitive data — or does your cloud provider? ☐1 ☐2 ☐3 ☐4 | ☐ CRITICAL CELL Does your model provider process your data in plaintext on their infra? ☐1 ☐2 ☐3 ☐4 | Do you have cryptographic controls over what agents can access and output? ☐1 ☐2 ☐3 ☐4 |
| **PILLAR 4 Operational Control** Do you have real-time visibility? | Real-time visibility into energy use and carbon intensity per workload? ☐1 ☐2 ☐3 ☐4 | Real-time visibility into compute utilization and cost across all GPUs? ☐1 ☐2 ☐3 ☐4 | Can you detect a residency violation or breach within minutes, not hours? ☐1 ☐2 ☐3 ☐4 | Real-time visibility into model behavior, quality, and decision provenance? ☐1 ☐2 ☐3 ☐4 | ☐ CRITICAL CELL Can you monitor, pause, or audit every agent action in real time? ☐1 ☐2 ☐3 ☐4 |
| **PILLAR 5 Contractual Control** Do you have enforceable rights? | Do energy contracts include audit rights, portability, and exit provisions? | Do compute agreements protect against unilateral capacity restrictions? | Do data center agreements include audit rights and deletion certification? | ☐ CRITICAL CELL Do you own your query logs and outputs — or does the provider retain? | ☐ CRITICAL CELL Do agent agreements cover liability for autonomous decisions and data exit? ☐1 ☐2 ☐3 ☐4 |

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

**THE GAP BETWEEN THE TWO SCORES IS YOUR PROGRAMME**

*The 0–160 score gives you your destination. The 5×5 matrix gives you your starting point. The distance between them — cell by cell — is the governance program.*

## Why This Assessment Matters for Private Equity

**92%**

of advanced AI chips from a single company in a geopolitically contested region

**70%**

of global AI compute controlled by five providers — all of which have ongoing access to your queries by design

**100%**

of PE firms using external AI models are sharing deal intelligence with provider infrastructure by design

AI governance gaps do not disappear at close — they transfer. Every portfolio company acquired with undisclosed AI infrastructure dependencies, unauditable model outputs, and provider agreements that predate applicable regulation creates post-acquisition liability that no representation and warranty policy was written to cover. The time to find those gaps is before you own them.

# 2 BACKGROUND: FIVE PILLARS & THE PRIVATE EQUITY AI CHALLENGE

The Five Pillars are the governance dimensions applied across all five AI ecosystems in the matrix. Understanding them ensures your deal team and investment committee are assessing the same governance dimensions at each cell.

## PILLAR 1 — Jurisdictional Control

*Where data and compute actually reside, and under which laws*

### What it means

- Ability to prove with technical evidence the physical and legal location of every AI workload
- Technical enforcement of data residency requirements — not just contractual promises
- Real-time visibility into cross-border data flows across all deal and portfolio data
- Automated prevention of unauthorized jurisdictional transfers of proprietary intelligence

### Why it matters for private equity

- SEC-registered investment advisers have record-keeping obligations that extend to AI system logs — records must be producible on examination demand
- Deal intelligence processed on foreign AI infrastructure may be subject to government demands the firm cannot challenge
- LP data sovereignty expectations are tightening — institutional LPs are beginning to require technical, not contractual, data residency assurances

### Common gaps in private equity AI

- Deal analysis queries processed on provider infrastructure in jurisdictions not mapped against SEC obligations
- Portfolio company financial data routed through provider systems in multiple jurisdictions without residency assessment
- LP information processed by external AI models under terms that do not address jurisdictional governance

## PILLAR 2 — Logical Control

*Who can access what, when, with immutable proof*

### What it means

- Identity-based access control enforced at infrastructure layer
- MFA for 100% of privileged accounts across deal and portfolio systems
- Just-in-time access with automatic expiration for sensitive deal environments
- Immutable audit logs of every authorization in firm-controlled systems
- Automated revocation upon departure of investment professionals

### Why it matters for private equity

- 68% of security incidents stem from excessive privileges — in PE, insider access to deal intelligence is the primary competitive risk
- SEC examination authority extends to records of who accessed AI systems processing deal and portfolio data
- Model provider engineers have technical access to deal queries and outputs by design — a logical control gap most PE firms have not addressed

### Common gaps in private equity AI

 

- Former investment professionals retain access to deal intelligence systems after departure
- Model provider support engineers can access deal query environments without firm-controlled logging
- Access reviews occur periodically but deal intelligence exposure is continuous

## PILLAR 3 — Technical Control

*Cryptographic sovereignty over deal intelligence, models, and compute*

### What it means

- Encryption keys held in firm-controlled HSMs (HYOK minimum, BYOK acceptable)
- Provider cannot decrypt deal intelligence without firm's active participation
- Confidential computing protecting proprietary data during AI model processing
- Cryptographic attestation of workload integrity across all deal systems

### Why it matters for private equity

- Provider-managed encryption means providers can decrypt deal intelligence on government demand — including demands from foreign governments the firm cannot challenge
- Without key control, a firm's ownership of its deal intelligence is a legal assertion unsupported by technical architecture
- BYOK reduces unauthorized access events by 40%; HYOK reduces incidents by 65% and provides the legal instrument of sovereignty over deal data

### Common gaps in private equity AI

- Standard AI model API processing of proprietary deal thesis and portfolio intelligence in plaintext
- BYOK implemented for cloud storage but not extending to AI model inference layer
- No confidential computing for high-sensitivity deal analysis AI workloads

## PILLAR 4 — Operational Control

*Real-time visibility into what is happening across the deal and portfolio AI estate*

### What it means

- Centralized telemetry from all infrastructure — including deal AI, portfolio monitoring, and LP reporting systems
- Real-time anomaly detection within minutes, not hours or days
- Continuous compliance monitoring across all AI systems touching deal and portfolio intelligence
- Immutable audit trails with cryptographic integrity verification in firm-controlled systems

### Why it matters for private equity

- 74% of enterprises lack end-to-end visibility into AI operational data flows — in PE, this includes deal pipeline AI and portfolio company monitoring systems
- Mean time to detect incidents is 127 hours under provider-managed infrastructure vs 34 hours with firm-controlled monitoring
- SEC examination increasingly focuses on AI in investment processes — retrospective log reconstruction is not a viable examination strategy

### Common gaps in private equity AI

- Deal AI logs collected in vendor systems but not actively monitored in firm-controlled infrastructure
- Portfolio company AI monitoring verified through vendor dashboards rather than firm-controlled telemetry
- Agent actions processing deal screening and LP reporting invisible to firm SIEMs

## PILLAR 5 — Contractual Control

*Enforceable legal rights over deal intelligence, AI outputs, and provider accountability*

 

### What it means

- Explicit audit rights over AI providers without provider consent requirement
- Subprocessor approval workflows with LP notification rights
- Deal intelligence portability in open formats with defined exit timelines
- Material breach definitions with enforceable remedies beyond nominal fee caps
- Liability alignment for proprietary intelligence exposure

### Why it matters for private equity

- Standard AI model API terms do not provide ownership of interaction logs, certified deletion of deal intelligence, or audit rights exercisable by SEC examiners or LP due diligence teams
- Most PE firm AI vendor agreements were not reviewed by fund counsel against SEC Advisers Act record-keeping obligations or ILPA governance standards
- LP due diligence questionnaires are beginning to ask specifically about AI governance — weak contractual terms create LP relationship risk that is already materializing

### Common gaps in private equity AI

- Model API agreements allow unlimited subprocessing of deal intelligence without firm notification
- No audit rights over AI model providers without provider consent
- No deletion verification for deal intelligence — only provider attestation
- Exit costs uncapped, creating dependency on providers who can change terms at any time

## The Private Equity AI Governance Challenge

Private equity firms manage some of the most competitively sensitive intelligence in the financial system — deal pipelines, proprietary investment theses, portfolio company operational data, and LP capital commitments. The AI systems being deployed across deal sourcing, due diligence, portfolio monitoring, and LP reporting are processing that intelligence on external infrastructure, under standard commercial terms, with logs held in vendor systems.

The challenge is not access to AI. Every GP has access. The challenge is governance. Who owns the deal intelligence AI processes. Who governs the models. Who ensures SEC compliance, LP transparency, and competitive intelligence protection — not through provider promises, but through technical architecture that enforces it.

> *The most acute governance gap for most PE firms is not a government legal demand. It is that external model providers — Anthropic, OpenAI, Microsoft, Google — have ongoing access to deal queries, investment thesis logic, and portfolio company intelligence by design. Every day. In every API call. That is the operational reality of how AI models are served.*

### The four PE AI use cases that create governance obligations

#### Deal Sourcing and Screening

AI systems scanning deal flow, scoring targets, and processing proprietary signal intelligence carry competitive intelligence obligations that standard API terms do not address. The deal thesis logic embedded in every AI screening query is processed on provider infrastructure.

#### Due Diligence Synthesis

AI-assisted data room ingestion, financial anomaly detection, and commercial diligence synthesis process confidential target company information and proprietary deal analysis under standard terms whose data governance provisions were not written for M&A confidentiality obligations.

#### Portfolio Company Monitoring

AI systems processing portfolio company financial data, operational metrics, and strategic intelligence for ongoing monitoring handle information whose confidentiality obligations extend to the portfolio company itself, its management, and the PE firm's LP base.

#### LP Reporting and Investor Relations

 

AI-assisted LP reporting and capital raising intelligence processes investor information whose governance obligations under the SEC Advisers Act, ILPA standards, and LP agreements most firms have not mapped to their AI vendor relationships.

 

# 3 THE ASSESSMENT

Complete all four sections honestly and comprehensively. Score what you can demonstrate with evidence — not what you believe is true or what providers have attested. Recommended team: Managing Partner, General Counsel, COO, CFO, Chief Compliance Officer, Head of Technology. Estimated time: 45–60 minutes.

> Score your current state, not your aspirations. A Level 1 score is accurate information. It is more valuable than an inflated score that does not reflect technical reality. The assessment is most useful when it reveals the true gap between where you are and where you need to be.

## SECTION 1
### REGULATORY & FIDUCIARY REQUIREMENTS

*How binding are your SEC, LP fiduciary, and fund governance obligations?*

**1.1 Which regulatory regimes govern your firm's AI operations?**

- ☐ **SEC-registered investment adviser with full Advisers Act obligations** *(10 pts)*
  *Record-keeping, fiduciary duty, and examination authority extend to all AI systems contributing to investment decisions and client communications.*
- ☐ **SEC-registered but with limited regulatory footprint** *(8 pts)*
  *Some Advisers Act obligations apply; LP fiduciary duties govern investment decision-making processes.*
- ☐ **Exempt reporting adviser or state-registered** *(6 pts)*
  *Reduced SEC examination exposure but LP fiduciary obligations and fund agreement governance requirements still apply.*
- ☐ **Non-US fund with EU or other jurisdictional obligations** *(3 pts)*
  *AIFMD, MiFID II, or equivalent frameworks apply; cross-border data governance complexity is significant.*
- ☐ **No formal regulatory registration** *(0 pts)*
  *Fund counsel oversight only; commercial best practices govern AI deployment.*

**1.1 Score Score: _____ / 10**

> SEC examination authority extends to AI systems contributing to investment decisions. AI system logs are records within the meaning of the Advisers Act. Most PE firms have not reviewed their AI vendor agreements against their record-keeping obligations.

**Notes:** _____

_____

**1.2 Can you demonstrate with technical evidence where every AI workload processing proprietary deal and LP data executes?**

- ☐ **Yes — with real-time technical verification** *(0 pts)*
  *Geo-fencing and continuous monitoring prove deal AI workload location at all times.*
- ☐ **Yes — based on contractual commitments** *(3 pts)*
  *Provider contractually commits to data residency; no independent technical verification.*
- ☐ **Partially — core systems verified; deal AI locations uncertain** *(6 pts)*
  *Fund administration tracked; AI deal screening and diligence synthesis locations unmapped.*
- ☐ **No — relying on provider attestation** *(9 pts)*
  *Provider claims compliance; no independent verification mechanism.*
- ☐ **Unknown** *(10 pts)*
  *No visibility into where AI processing of deal and LP data executes.*

**1.2 Score Score: _____ / 10**

> Red Flag: If the SEC served an examination demand today, could you produce a jurisdictional map of every AI system processing deal intelligence within 24 hours?

**Notes:** _____

_____

 

### 1.3 Could you produce a complete AI decision audit trail for any deal analysis or investment recommendation from 18 months ago within 24 hours?

- □ **Yes — automated export from firm-controlled systems** *(0 pts)*
  *Immutable logs covering deal screening, diligence AI, portfolio monitoring, and LP reporting.*
- □ **Yes — but requires manual compilation** *(4 pts)*
  *Logs exist but scattered across vendor systems; significant reconstruction effort required.*
- □ **Partially — core records available but AI model interaction logs absent** *(7 pts)*
  *Fund records available; model provider interaction logs not held in firm systems.*
- □ **No — AI system logs incomplete or held only by providers** *(10 pts)*
  *Direct SEC examination vulnerability: deal intelligence audit trails do not exist in firm-controlled form.*

**1.3 Score Score: _____ / 10**

Advisers Act Rule 204-2 requires registered investment advisers to maintain records of investment decisions and the basis for recommendations. AI system logs that contribute to those decisions are within the scope of this requirement.

**Notes:** _____

_____

### 1.4 Does your AI infrastructure involve proprietary intelligence that creates material risk if disclosed to competitors, foreign governments, or unauthorized parties?

- □ **Yes — active deal pipeline, investment thesis logic, and target company analysis** *(10 pts)*
  *Disclosure of deal pipeline would compromise negotiating position, LP trust, and competitive advantage simultaneously.*
- □ **Yes — portfolio company financial data and strategic intelligence** *(8 pts)*
  *Unauthorized disclosure of portfolio company data would breach confidentiality obligations to management teams and co-investors.*
- □ **Yes — LP capital commitment data and investor information** *(7 pts)*
  *LP information disclosed without consent creates Advisers Act violations and LP relationship damage.*
- □ **Some sensitivity — fund performance data and operational metrics** *(4 pts)*
  *Moderate competitive sensitivity; not mission-critical confidentiality.*
- □ **No material proprietary intelligence risk** *(0 pts)*
  *Non-sensitive AI workloads only.*

**1.4 Score Score: _____ / 10**

**Notes:** _____

_____

## SECTION 1 TOTAL: _____ / 40

| Score | Interpretation |
|---|---|
| **0–10** | Low regulatory pressure; commercial cloud with standard governance sufficient. Annual compliance reviews adequate. |
| **11–20** | Moderate requirements; enhanced contractual controls required — BYOK encryption, explicit audit rights, subprocessor transparency. |
| **21–30** | High obligations; sovereignty architecture essential. HYOK minimum for deal intelligence workloads; continuous monitoring mandatory. |
| **31–40** | Critical regulatory obligations; SEC-grade AI governance required. Fund counsel must review all AI vendor agreements immediately. |

 

# SECTION 2
## STRATEGIC IMPORTANCE OF AI

*How central is AI to deal sourcing, value creation, and competitive positioning?*

### 2.1 How critical is AI to your firm's deal sourcing and investment decision-making?

- □ **AI is existential — deal flow, screening, and portfolio monitoring cannot operate without it** *(10 pts)*
  *AI drives the majority of deal origination, target identification, and portfolio performance analysis.*
- □ **AI is strategic — core to competitive positioning and LP differentiation** *(8 pts)*
  *AI-driven deal sourcing and portfolio intelligence are primary competitive advantages marketed to LPs.*
- □ **AI is important — significant operational value across the deal cycle** *(5 pts)*
  *AI improves deal screening efficiency and portfolio monitoring but manual fallbacks exist.*
- □ **AI is beneficial — productivity enhancement for deal teams** *(2 pts)*
  *AI assists deal professionals but is not embedded in core investment decision workflows.*
- □ **AI is experimental — limited deployment** *(0 pts)*
  *Pilot projects only; no material investment process impact if ceased.*

**2.1 Score Score: _____ / 10**

**Notes:** _____

_____

### 2.2 What percentage of your firm's critical investment decisions will be AI-influenced within 3 years?

- □ **Greater than 75%** *(10 pts)*
  *AI will drive deal sourcing, screening, diligence synthesis, and portfolio monitoring at scale.*
- □ **50–75%** *(8 pts)*
  *AI drives most deal origination and portfolio functions; AI outages create significant operational disruption.*
- □ **25–50%** *(5 pts)*
  *AI supports significant functions but human judgment remains primary.*
- □ **10–25%** *(2 pts)*
  *AI supplements traditional deal-making in limited areas.*
- □ **Less than 10%** *(0 pts)*
  *AI is peripheral to investment process.*

**2.2 Score Score: _____ / 10**

**Notes:** _____

_____

### 2.3 If your primary AI provider restricted access for 90 days, what would be the impact on the firm?

- □ **Existential — deal sourcing and portfolio monitoring materially degrade** *(10 pts)*
  *Core investment functions disrupted; LP reporting delayed; competitive disadvantage materializes immediately.*
- □ **Severe — major disruption to deal pipeline and portfolio visibility** *(8 pts)*
  *Operations degraded 50–80%; LP relationship strain; deal execution delays.*
- □ **Significant — meaningful degradation requiring manual workarounds** *(5 pts)*
  *Operations degraded 20–50%; deal team capacity strained.*
- □ **Manageable — inconvenient but investment operations continue** *(2 pts)*
  *Operations degraded 10–20%; traditional processes viable.*
- □ **Minimal — deal operations continue normally** *(0 pts)*
  *AI is supplementary; no material investment process impact.*

**2.3 Score Score: _____ / 10**

**Notes:** _____

_____

### 2.4 Do you train proprietary AI models that represent strategic firm intelligence?

- □ **Yes — proprietary models are core IP and competitive moat** *(10 pts)*
  *Custom deal sourcing models, target scoring algorithms, and portfolio intelligence engines valued as primary competitive assets.*

 

- □ **Yes — custom models provide significant competitive advantage** *(7 pts)*
  *Fine-tuned models with measurable sourcing and diligence advantage.*
- □ **Some proprietary models, primarily commercial** *(4 pts)*
  *Hybrid approach with proprietary fine-tuning for deal-specific applications.*
- □ **Primarily commercial with deal-specific prompt engineering** *(2 pts)*
  *Light customization for investment context.*
- □ **No — commercial AI models used as-is** *(0 pts)*
  *Standard API calls; no proprietary training.*

**2.4 Score Score: _____ / 10**

**Notes:** _____
_____

## SECTION 2 TOTAL: _____ / 40

 

# SECTION 3
## RISK TOLERANCE & SOVEREIGNTY REQUIREMENTS

*What level of third-party dependency on deal and LP intelligence is acceptable?*

### 3.1 Your firm's risk tolerance for AI infrastructure dependency on deal intelligence:

- □ **Zero tolerance — all deal intelligence AI must be on controlled infrastructure** *(10 pts)*
  *Large-cap GP with institutional LPs requiring demonstrable data sovereignty and SEC examination readiness.*
- □ **Very low — dependency must be minimized and contractually controlled** *(7 pts)*
  *Mid-market GP where LP due diligence and SEC registration make strong controls essential.*
- □ **Moderate — acceptable if governance frameworks exist** *(4 pts)*
  *Emerging manager where operational efficiency is balanced against competitive intelligence risk.*
- □ **Higher — efficiency outweighs dependency concerns** *(1 pts)*
  *Early-stage fund where capital constraints limit governance investment.*

**3.1 Score Score: _____ / 10**

**Notes:** _____

_____

### 3.2 Can your AI providers currently decrypt the proprietary deal intelligence processed by your AI systems?

- □ **Yes — provider-managed encryption for deal AI workloads** *(10 pts)*
  *Provider can decrypt deal pipeline intelligence, investment theses, and portfolio company data at any time.*
- □ **Partially — BYOK for some workloads** *(6 pts)*
  *Better protection but provider can still access deal data during AI model processing (plaintext inference).*
- □ **No — HYOK with confidential computing for all deal intelligence** *(0 pts)*
  *Provider cannot access deal intelligence during processing. Meets highest competitive intelligence protection standard.*
- □ **Unknown — encryption key custody for AI workloads not mapped** *(10 pts)*
  *Requires immediate fund counsel attention. A direct competitive intelligence exposure.*

**3.2 Score Score: _____ / 10**

Under HYOK with confidential computing, a government demand on your AI provider cannot result in disclosure of proprietary deal intelligence without your knowledge. Under standard API terms, disclosure may occur without your awareness.

**Notes:** _____

_____

### 3.3 If a government served your AI model provider with a legal demand for your deal intelligence and LP data, would you:

- □ **Learn about it only after disclosure** *(10 pts)*
  *Deal pipeline and LP information disclosed without firm knowledge — direct competitive intelligence and fiduciary breach.*
- □ **Be notified but unable to prevent disclosure** *(7 pts)*
  *May learn within 24 hours but practical control over deal intelligence disclosure is limited.*
- □ **Have legal right to challenge before disclosure** *(4 pts)*
  *Can seek judicial review; outcome uncertain but process meaningful.*
- □ **Have technical controls preventing disclosure without firm participation** *(0 pts)*
  *HYOK plus confidential computing. Full competitive intelligence sovereignty.*

**3.3 Score Score: _____ / 10**

**Notes:** _____

_____

### 3.4 Does your LP base require demonstrable AI sovereignty in your fund operations?

- □ **Absolutely — LP due diligence and side letter requirements demand technical sovereignty** *(10 pts)*
  *Institutional LPs explicitly requiring AI governance documentation; side letters addressing data protection.*
- □ **Yes — LPs expect meaningful governance control over deal and portfolio data** *(7 pts)*
  *Standard LP due diligence questionnaires now include AI governance questions.*

 

- □ **Somewhat — sovereignty enhances LP confidence but is not yet formally required** *(4 pts)*
  *Governance improvement is beneficial; formal LP requirements not yet imposed.*
- □ **No — LPs have not yet raised AI governance in due diligence** *(0 pts)*
  *Smaller LPs; AI governance not yet a selection criterion.*

**3.4 Score Score: _____ / 10**

**Notes:** _____
_____

## SECTION 3 TOTAL: _____ / 40

 

# SECTION 4
## FINANCIAL & OPERATIONAL CAPACITY
*Can the firm fund and operate sovereign AI infrastructure?*

### 4.1 Your firm's AI infrastructure capital budget over the next 3 years:

- □ **Greater than $500M — can build dedicated sovereign infrastructure** *(10 pts)*
  *Largest GPs with direct GPU procurement; multi-strategy sovereign infrastructure.*
- □ **$100M–$500M — can build hybrid/regional sovereign infrastructure** *(7 pts)*
  *Large buyout firm with selective GPU ownership; sovereign cloud partnerships.*
- □ **$20M–$100M — can enhance governance with limited infrastructure** *(4 pts)*
  *Enhanced contractual terms; BYOK/HYOK; dedicated compliance and technology function.*
- □ **Less than $20M — rental with governance overlay only** *(1 pts)*
  *Commercial cloud with enhanced contractual terms; small governance team.*

**4.1 Score Score: _____ / 10**

**Notes:** _____
_____

### 4.2 Do you have in-house capability to operate AI infrastructure at institutional scale?

- □ **Yes — deep AI infrastructure expertise in-house** *(0 pts)*
  *Technology team with hyperscale experience; GPU cluster management; 24/7 operations.*
- □ **Partially — strong technology team; AI governance building** *(4 pts)*
  *Traditional fund technology team; AI governance programme developing.*
- □ **No — would require significant hiring** *(8 pts)*
  *Primarily application-focused teams; 20–100 specialized FTEs required.*
- □ **No — prefer outsourced managed services** *(10 pts)*
  *Focus on investment operations; limited appetite for AI infrastructure complexity.*

**4.2 Score Score: _____ / 10**

**Notes:** _____
_____

### 4.3 What is your acceptable timeline for achieving meaningful AI sovereignty for deal intelligence?

- □ **Already achieved or less than 6 months** *(0 pts)*
  *Infrastructure owned or contracted; governance implemented; SEC examination ready.*
- □ **6–12 months — LP pressure or regulatory urgency is imminent** *(3 pts)*
  *Upcoming fund raise or SEC examination cycle driving urgency.*
- □ **12–24 months — strategic initiative with business case** *(6 pts)*
  *Documented competitive intelligence risk driving phased investment.*
- □ **24–36 months — long-term vision** *(9 pts)*
  *Building optimal solution with lower execution risk.*
- □ **Greater than 36 months or no urgency** *(10 pts)*
  *AI governance deferred; reassess as AI adoption grows and LP expectations develop.*

**4.3 Score Score: _____ / 10**

**Notes:** _____
_____

### 4.4 Annual AI compute spending across fund operations (current or projected within 2 years):

- □ **Greater than $50M/year** *(10 pts)*
  *Break-even on sovereign infrastructure: 5–7 years. Profile: largest global PE platforms.*
- □ **$20M–$50M/year** *(7 pts)*
  *Large buyout or multi-strategy firm with significant AI deployment.*
- □ **$5M–$20M/year** *(4 pts)*
  *Build not economical at this scale; governance overlay most appropriate.*
- □ **Less than $5M/year** *(1 pts)*
  *Focus investment on contractual controls and governance framework.*

 

**4.4 Score Score: _____ / 10**

**Notes:** _____

_____

## SECTION 4 TOTAL: _____ / 40

 

# 4 TOTAL ASSESSMENT SCORE & INTERPRETATION

| Section | Max | Your Score |
|---|---|---|
| Section 1 — Regulatory & Fiduciary Requirements | 40 | _____ |
| Section 2 — Strategic Importance of AI | 40 | _____ |
| Section 3 — Risk Tolerance & Sovereignty | 40 | _____ |
| Section 4 — Financial & Operational Capacity | 40 | _____ |

## TOTAL STRATEGIC ASSESSMENT SCORE: _____ / 160

| Score | Strategy | What it means for your firm |
|---|---|---|
| 0–40 | RENT | AI is supplementary. Managed cloud with enhanced contractual controls. Immediate priority: fund counsel review of all AI vendor agreements. |
| 41–80 | RENT + GOVERN | AI is important. Enhanced contractual controls and BYOK encryption required. Models column must be contractual priority before any AI capability expansion. |
| 81–120 | COMPOSE | AI is strategic. Hybrid sovereign architecture needed. Sovereign infrastructure for deal intelligence and LP data. Managed cloud for less-sensitive workloads. |
| 121–160 | BUILD | AI is existential. Full sovereign infrastructure required. Complete five-pillar control framework across every ecosystem. |

 

# 5 STEP 1 — THE 5×5 CONTROL MATRIX

**Complete this section first. Score what you can demonstrate with evidence.**

Work through all 25 cells — one ecosystem at a time, all five pillars for each. Score each cell 1 to 4. Use the notes field to record the evidence basis for your score. Complete this before reviewing benchmarks or completing the strategic assessment.

> Score what you can demonstrate — not what you believe, not what contracts assert, not what providers have attested. A Level 1 score is honest information. It is more valuable than an inflated score that does not survive SEC examination.

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? □1 □2 □3 □4 | Under which jurisdiction does GPU compute actually execute? □1 □2 □3 □4 | Can you prove where every data center workload executes? □1 □2 □3 □4 | Which jurisdiction governs model training, storage, and serving — and does it grant access to weights? □1 □2 □3 □4 | Where do agent actions execute and where do decision logs reside? □1 □2 □3 □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? □1 □2 □3 □4 | Who accesses GPU clusters — including provider support engineers? □1 □2 □3 □4 | Can you prove no unauthorized access in the past 18 months? □1 □2 □3 □4 | □ CRITICAL CELL Can your model provider access your queries and outputs — in your logs? □1 □2 □3 □4 | Who can approve or halt agents — are all actions logged in your systems? □1 □2 □3 □4 |
| **PILLAR 3 Technical Control** Who holds the encryption keys? | Do you control encryption keys for energy and ESG data? □1 □2 □3 □4 | Who controls keys for AI training data and model weights? □1 □2 □3 □4 | Do you hold HYOK for all sensitive data — or does your cloud provider? □1 □2 □3 □4 | □ CRITICAL CELL Does your model provider process your data in plaintext on their infra? □1 □2 □3 □4 | Do you have cryptographic controls over what agents can access and output? □1 □2 □3 □4 |
| **PILLAR 4 Operational Control** Do you have real-time visibility? | Real-time visibility into energy use and carbon intensity per workload? □1 □2 □3 □4 | Real-time visibility into compute utilization and cost across all GPUs? □1 □2 □3 □4 | Can you detect a residency violation or breach within minutes, not hours? □1 □2 □3 □4 | Real-time visibility into model behavior, quality, and decision provenance? □1 □2 □3 □4 | □ CRITICAL CELL Can you monitor, pause, or audit every agent action in real time? □1 □2 □3 □4 |
| **PILLAR 5 Contractual Control** Do you have enforceable rights? | Do energy contracts include audit rights, portability, and exit provisions? | Do compute agreements protect against unilateral capacity restrictions? | Do data center agreements include audit rights and deletion certification? | □ CRITICAL CELL Do you own your query logs and outputs — or does the provider retain? | □ CRITICAL CELL Do agent agreements cover liability for autonomous decisions and data exit? □1 □2 □3 □4 |

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

## ECOSYSTEM 1 — POWER | *Energy infrastructure*

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

**Where is the energy data powering your AI deal sourcing, portfolio monitoring, and fund administration platforms processed — and under which jurisdiction?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*

**Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 AI investment platforms — and is that access immutably logged in systems producible for SEC examination within 24 hours?**

**Level 1** □ *Reactive — relying on provider assurances; no independent verification*
**Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
**Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
**Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 across fund operations?**

**Level 1** □ *Reactive — relying on provider assurances; no independent verification*
**Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
**Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
**Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 LP sustainability requirements?**

**Level 1** □ *Reactive — relying on provider assurances; no independent verification*
**Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
**Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
**Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 SEC fiduciary obligations and the service provider oversight requirements LPs must demonstrate?**

**Level 1** □ *Reactive — relying on provider assurances; no independent verification*
**Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
**Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
**Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 deal screening, due diligence synthesis, and portfolio company analysis actually execute?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 proprietary deal intelligence and LP data — including cloud provider support engineers — and is that access logged in your own systems?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 datasets, and proprietary deal intelligence underlying your investment AI systems?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 deal sourcing, due diligence, and portfolio company monitoring?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**P4×E2 Score:** _____ / 4

**Notes:** _____

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

**Do your compute and cloud agreements protect proprietary deal and LP data sovereignty and include workload portability rights sufficient for migration without deal intelligence exposure?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*

 

**Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
**Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
**Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 deal record, LP communication, portfolio company data, and AI-generated investment analysis executes — satisfying SEC examination and LP fiduciary oversight requirements?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 deal record, LP information, portfolio company data, or AI-generated investment analysis in the past 18 months — satisfying SEC examination and LP oversight requirements?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**P2×E3 Score:** _____ / 4

**Notes:** _____

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

**Do you hold your own encryption keys (HYOK) for all proprietary deal data — pipeline intelligence, portfolio company financials, LP information, investment theses — or does your cloud provider manage and control them?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**P3×E3 Score:** _____ / 4

**Notes:** _____

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

**Can you detect an unauthorized access to deal intelligence, a portfolio AI anomaly, or a data residency violation within minutes — satisfying SEC fiduciary obligations and LP oversight requirements?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 LP notification rights, certified deletion of proprietary data, and exit terms consistent with SEC record retention requirements?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 proprietary deal intelligence, investment thesis data, portfolio company financials, and LP capital commitment information trained, stored, and served?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**P1×E4 Score:** _____ / 4

**Notes:** _____

_____

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

**Can your AI model providers access the proprietary data your models process — deal pipeline intelligence, investment theses, portfolio company financials, LP commitments — and is that access logged in your own infrastructure?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**P2×E4 Score:** _____ / 4

**Notes:** _____

_____

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

**When your AI models process proprietary deal intelligence, portfolio company data, or LP information, does the model provider process this data in plaintext on their own infrastructure?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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, deal analysis quality, and portfolio monitoring accuracy across all AI systems contributing to investment decisions and fund management?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 deal intelligence processed, diligence logs, and investment outputs — or does the provider retain residual rights over your most competitively sensitive information?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 deal sourcing, due diligence, and portfolio monitoring agents execute — and where do their action logs and deal interaction records legally reside?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**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 deal sourcing, due diligence, and portfolio monitoring agents — and is every agent action immutably logged in institution-controlled systems accessible to GPs and regulators?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**P2×E5 Score:** _____ / 4

**Notes:** _____

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

**Do you have cryptographic enforcement over what your autonomous deal sourcing, diligence, and portfolio monitoring agents can access, modify, and transmit — protecting competitive intelligence beyond policy rules?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**P3×E5 Score:** _____ / 4

**Notes:** _____

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

**Can you monitor, pause, or audit every autonomous agent processing deal sourcing, due diligence, portfolio monitoring, and LP reporting — in real time, with complete action logs accessible to GPs and regulators?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**P4×E5 Score:** _____ / 4

**Notes:** _____

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

 

**Do your agentic AI vendor agreements include firm liability for autonomous deal and portfolio decisions, audit rights exercisable by LPs and SEC examiners, and data exit provisions covering all proprietary data processed by agents?**

- **Level 1** □ *Reactive — relying on provider assurances; no independent verification*
- **Level 2** □ *Evolving — contractual controls only; limited technical enforcement*
- **Level 3** □ *Governed — active monitoring, contractual rights, partial technical controls*
- **Level 4** □ *Sovereign — cryptographically verifiable; auditable by SEC and LP due diligence*

**P5×E5 Score: _____ / 4**

Notes: _____

_____

## ECOSYSTEM 5 — AGENTS SUBTOTAL: _____ / 20

## Matrix Scoring Summary

| 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 |

## YOUR MATRIX SCORE: _____ / 100 *Record this before proceeding to Step 2.*

**Now move to Section 6 — Sector Benchmarks — to understand what your score means relative to your peers before reviewing recommendations.**

 

# 8 CRITICAL CELL ANALYSIS

A critical cell is any matrix intersection at Level 1 where the pillar-ecosystem combination creates material competitive intelligence, regulatory, or LP relationship risk. These six intersections represent the highest-risk critical cells for private equity firms — based on the structural nature of AI adoption in the deal cycle and the gap between standard API terms and PE governance obligations.

## Logical × Models

**Typical exposure:** *Almost universal at Level 1*

**Why it matters:** External AI model providers process proprietary deal intelligence — pipeline data, investment thesis logic, target company analysis — by design in every API call. Interaction logs are in the provider's systems. When an LP asks who had access to their fund's deal intelligence on a specific date, or when the SEC examines AI's role in investment decisions, the complete answer requires logs the firm does not hold.

**Minimum action:** Negotiate enterprise model API agreements with explicit deal intelligence access logging in firm-controlled systems. Build model provider interaction logging into your SIEM before expanding any AI capability in the deal cycle.

## Technical × Models

**Typical exposure:** *Most PE firms at Level 1–2*

**Why it matters:** Proprietary deal intelligence — the source material of competitive advantage in PE — is processed in plaintext by external AI providers under standard API terms. BYOK protects data at rest but not during model inference. Every deal query is a technical exposure of the firm's most sensitive competitive intelligence to provider infrastructure.

**Minimum action:** Evaluate confidential computing options for all model inference involving deal intelligence. For the most sensitive workloads — active deal pipeline, investment thesis data — treat confidential computing as the minimum acceptable technical control before expanding AI use in the deal cycle.

## Contractual × Models

**Typical exposure:** *Most PE firms at Level 1–2*

**Why it matters:** Standard AI model API terms do not provide: explicit ownership of deal interaction logs; certified deletion of deal intelligence submitted as queries; audit rights without provider consent; or liability alignment for competitive intelligence exposure. Most PE fund counsel have reviewed cloud infrastructure contracts but have not applied the same standard to AI model API agreements.

**Minimum action:** Fund counsel must review every AI model provider agreement against SEC Advisers Act record-keeping obligations, LP fiduciary duties, and the competitive intelligence protection obligations the firm owes to its portfolio companies and co-investors. Negotiate custom enterprise terms before expanding AI use in deal-sensitive workflows.

## Contractual × Agents

**Typical exposure:** *Almost universal at Level 1*

**Why it matters:** Autonomous agents operating in deal sourcing, portfolio monitoring, and LP reporting workflows are governed by startup-level vendor agreements that do not address firm liability for autonomous deal intelligence decisions. As agents become embedded in the deal cycle — processing pipeline data, generating investment summaries, communicating with portfolio companies — the contractual framework governing liability for errors and data breaches is undefined.

**Minimum action:** Do not deploy agentic AI in deal-sensitive workflows without custom contractual terms: liability for autonomous deal decisions; audit rights over every agent action; data exit provisions for agent-processed deal intelligence; and human-in-the-loop requirements for actions above defined deal significance thresholds.

## Operational × Agents

**Typical exposure:** *Almost universal at Level 1*

**Why it matters:** Agent actions, deal queries, portfolio data accesses, and LP communication outputs are largely invisible to firm-controlled SIEMs. Most agentic platforms log to their own infrastructure. Without real-time agent telemetry in firm-controlled monitoring systems, SEC examination readiness and LP oversight of autonomous deal intelligence is structurally impossible.

 

**Minimum action:** Establish minimum agent telemetry requirements before any production deployment: every agent action, tool call, deal data access, and output must be logged to firm-controlled systems in real time, in a format accessible to SEC examiners and LP due diligence teams upon request.

## Jurisdictional × Models

**Typical exposure:** *Most PE firms at Level 1–2*

**Why it matters:** AI models processing proprietary deal intelligence may be trained, stored, and served from jurisdictions whose laws create governance conflicts with the firm's competitive intelligence obligations and LP data protection commitments. A government demand served on a model provider in its home jurisdiction may result in disclosure of deal pipeline data without the firm's knowledge or ability to challenge.

**Minimum action:** Map the full jurisdictional footprint of every AI model provider relationship — training location, serving infrastructure, interaction log residency. Assess alignment against SEC record-keeping obligations and LP data protection expectations. For deal intelligence with specific residency requirements, evaluate jurisdiction-aligned model alternatives.

 

# 11 RED FLAGS & GREEN FLAGS

## Red Flags — When Not to Build Sovereign Infrastructure

### Deal AI strategy is still experimental

**Warning signs:** *No production AI workloads in the deal cycle yet; use cases being validated; ROI unproven.*

**Alternative:** Build governance framework and contractual controls first. Achieve Level 2 across the matrix before committing capital to infrastructure.

### No AI infrastructure expertise in-house

**Warning signs:** *No team members with hyperscale experience; underestimating operational complexity.*

**Alternative:** Build with partner support. Rent with enhanced governance while building capability, or partner with a managed sovereign cloud provider.

### Fund is in market with compressed timeline

**Warning signs:** *Fundraising timeline incompatible with 18–36 month sovereignty programme; LP commitments at risk.*

**Alternative:** Focus on contractual governance and LP due diligence documentation first. Sovereignty programme begins post-close; contractual controls provide interim protection.

### Portfolio company AI complexity dominates

**Warning signs:** *Multiple portfolio companies with disparate AI deployments create governance complexity that exceeds fund-level programme capacity.*

**Alternative:** Sequence fund-level governance first. Portfolio company AI governance becomes a value creation initiative post-investment, not a simultaneous fund-level programme.

### Partnership lacks commitment or understanding

**Warning signs:** *Treating AI governance as a technology project rather than a competitive intelligence and fiduciary imperative.*

**Alternative:** Build the LP due diligence case first. LP pressure is the most effective driver of partnership commitment to AI governance investment.

## Green Flags — When to Accelerate

### Institutional LP conducting formal AI governance due diligence

**Indicator:** *LP due diligence questionnaire now includes specific AI governance questions; LP making commitment contingent on governance documentation.*

**Action:** **Accelerate. The LP due diligence conversation is already happening. Lead it with a scored matrix rather than respond to questions without one.**

### SEC examination has included AI questions

**Indicator:** *SEC examination has asked about AI's role in investment decisions and the records the firm maintains.*

**Action:** **Accelerate. Build the evidence package before the next examination cycle begins.**

### Competitor has announced sovereign deal AI infrastructure

**Indicator:** *A competing GP has publicly positioned AI sovereignty governance as a fundraising differentiator.*

**Action:** **Governance is entering the LP selection conversation. Build it now rather than respond to it in an LP meeting.**

### Portfolio company AI governance gap identified post-close

**Indicator:** *M&A due diligence failed to identify AI governance gaps that are now creating post-acquisition liability.*

**Action:** **The M&A due diligence AI governance product is immediately relevant. Build it before the next deal.**

 

### Annual AI compute spend approaching $10M

**Indicator:** *AI operational cost reaching the threshold where governance investment becomes economically rational.*

**Action:** **Conduct a detailed TCO analysis. The governance programme ROI becomes compelling at this spend level.**

### New fund raise in 12–24 months

**Indicator:** *Next flagship fund raise will face LP due diligence questions about AI governance that the current posture cannot answer.*

**Action:** **The fund raise timeline is the deadline. Start the programme now to present a credible AI governance story at first LP close.**

 

# GLOSSARY OF KEY TERMS

The following definitions cover the principal terms used in this Private Equity Edition of the AI Sovereignty Assessment.

**5×5 Control Matrix** — The current-state governance diagnostic. Applies five control pillars independently to each of five AI ecosystems, producing 25 scored governance intersections and a total of 0–100. Reveals where governance is technically enforced and where it relies on provider assurances.

**0–160 Strategic Assessment** — The forward-looking strategic decision tool. Evaluates regulatory obligations, AI dependency, risk tolerance, and financial capacity to determine the appropriate AI infrastructure strategy: Rent, Rent + Govern, Compose, or Build.

**Advisers Act Rule 204-2** — SEC record-keeping requirement for registered investment advisers. Requires maintenance of records related to investment decisions and recommendations — including AI system logs that contribute to those decisions.

**Agentic AI** — AI systems that autonomously plan, decide, and execute multi-step tasks. In PE, agents process deal sourcing, pipeline screening, portfolio monitoring, and LP reporting functions — the least-governed layer of most PE AI deployments.

**BYOK (Bring Your Own Key)** — An encryption model in which the firm generates and manages its own encryption keys. Meets standard compliance requirements but does not prevent provider access during AI model inference — deal intelligence is still processed in plaintext.

**Confidential Computing** — Hardware-based technology protecting data while being processed — preventing providers from accessing deal intelligence in plaintext during model inference. Available through Intel TDX, AMD SEV, and NVIDIA Confidential Computing.

**Critical Cell** — Any matrix intersection at Level 1 where the pillar-ecosystem combination creates material competitive intelligence, regulatory, or LP relationship risk.

**Deal Intelligence** — The proprietary competitive advantage embedded in a PE firm's AI systems: active pipeline data, investment thesis logic, target company analysis, portfolio company financials, and LP capital commitment information.

**HYOK (Hold Your Own Key)** — An encryption model in which the firm retains full custody of encryption keys in its own HSM. The provider cannot decrypt deal intelligence under any circumstance — including government demands served on the provider.

**ILPA (Institutional Limited Partners Association)** — Industry body representing institutional LPs. ILPA standards and reporting templates are evolving to include AI governance due diligence requirements for PE fund managers.

**LP Due Diligence AI Governance** — The emerging standard by which institutional LPs — pension funds, sovereign wealth funds, endowments — assess PE GP AI governance as part of fund commitment due diligence. Increasingly includes matrix-level governance documentation.

**OLTAIX™** — Institutional AI's proprietary Control Tower — the Sovereign Intelligence Plane that orchestrates and governs all five AI ecosystems in real time, including deal intelligence, portfolio monitoring, and LP reporting AI.

**SEC Advisers Act** — Investment Advisers Act of 1940. Governs SEC-registered investment advisers including PE firms above AUM thresholds. Record-keeping, fiduciary duty, and examination authority provisions extend to AI systems contributing to investment decisions.

**Sovereign AI™** — The condition in which a PE firm's AI is fully owned, governed, and trusted — with every deal intelligence AI action traceable, every competitive intelligence workload technically protected, and every outcome accountable to GPs and LPs.

**The Institutional AI Stack™** — Institutional AI's proprietary architecture connecting all five AI ecosystems under one governed structure — designed to ensure deal intelligence AI receives the same competitive protection and regulatory governance as human investment processes.

 

## L E G A L  D I S C L A I M E R

*This document is provided solely for informational and educational purposes. It does not constitute legal, regulatory, investment, or other professional advice, and does not create an attorney-client or advisory relationship. Institutions should seek independent legal, regulatory, and technical counsel — including qualified securities counsel — before making decisions related to AI infrastructure, governance, or compliance.*

*All assessments, scores, and recommendations contained herein are illustrative and intended to support internal discussion and strategic planning. They do not represent an official certification, audit, or regulatory determination. Institutional AI assumes no responsibility or liability for any decisions or outcomes resulting from reliance on this material.*

*Nothing in this document constitutes legal advice regarding the SEC Advisers Act, LP fiduciary obligations, ILPA standards, or any other applicable law or regulation. Firms should consult qualified securities and fund counsel before making decisions based on this material.*