

FOR THE BOARD | EXECUTIVE BRIEF |
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.

<p>INSTRUMENT 1</p> <p>The 0–160 Strategic Assessment</p> <p><i>Sections 1–4 of this document</i></p> <p>Answers: What should our AI infrastructure strategy be?</p> <p>This is a forward-looking strategic decision tool. It evaluates your regulatory environment, how important AI is to your institutional mission, your risk tolerance, and your financial capacity. The output maps to one of four strategies:</p> <ul style="list-style-type: none"> ■ RENT (score 0–40) ■ RENT + GOVERN (score 41–80) ■ COMPOSE (score 81–120) ■ BUILD (score 121–160) <p><i>This tells you where you need to go.</i></p>	<p>INSTRUMENT 2</p> <p>The 5×5 Control Matrix</p> <p><i>Section 6 of this document</i></p> <p>Answers: Where is our governance today?</p> <p>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.</p> <p>Each of the 25 cells is scored 1 (Reactive) to 4 (Sovereign). Maximum total: 100 points.</p> <p><i>This tells you where you are standing today.</i></p>
--	--

THE RELATIONSHIP BETWEEN THE TWO SCORES
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.

A CONCRETE EXAMPLE

An institution completes the 0–160 assessment and scores 95. That falls in the 81–120 band — Compose is the right strategy. They need hybrid sovereign architecture with a protected core for sensitive data and proprietary systems.

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 are at Level 1.

The gap between 'you need to Compose' and 'you currently govern at Reactive level' is the entire work programme. The matrix identifies which of the 25 cells are at Level 1 — those are your priorities. In most institutions, the Models and Agents columns score lowest, meaning sensitive client and institutional data is being processed without adequate technical controls. Those Level 1 cells are the programme — not theoretical risk, but specific governance gaps with specific remediation steps.

THE FOUR ASSESSMENT DIMENSIONS		
Section 1	Regulatory & Fiduciary Requirements <i>How binding are your compliance and governance obligations?</i>	40 pts
Section 2	Strategic Importance of AI <i>How existential is AI to your institutional mission?</i>	40 pts
Section 3	Risk Tolerance & Sovereignty <i>What level of third-party AI dependency is acceptable?</i>	40 pts

Section 4	Financial & Operational Capacity <i>Can you fund and operate sovereign AI infrastructure?</i>	40 pts
------------------	---	---------------

WHY THIS MATTERS NOW

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 institutional sovereignty to the infrastructure providers that operate it.

FOUR OUTCOME STRATEGIES

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

FIVE QUESTIONS FOR THE BOARD

1. Can management prove — with technical evidence — where every AI workload processing sensitive institutional and client data executes today?
2. If our primary AI provider restricted access for 90 days, what would operationally happen to our core business functions?
3. Could we produce a complete AI decision audit trail from 18 months ago within 24 hours — for a regulator, a client, or a litigation proceeding?
4. Do our AI governance frameworks meet the standard our regulators, fiduciary duties, and client obligations require?
5. Do our model provider agreements give us explicit rights over the sensitive data processed, or does the provider retain residual rights over our most valuable institutional information?

Completion: 45–60 min | Score: 0–160 | information@institutional.ai | www.institutional.ai