



AI is a given. Control is not.

INSTITUTIONAL AI

OUR FRAMEWORK

Beyond the Standard

How the 5×5 Control Matrix™ relates to NIST, ISO 42001, and the EU AI Act — and why proving control is a different exercise than conforming to a framework.

The leading AI standards — the NIST AI Risk Management Framework, ISO/IEC 42001, the EU AI Act — define what good AI governance should look like. They are valuable, they are necessary, and we build on them. But conforming to a framework on paper is not the same as being able to prove control when a regulator, a board, or a fiduciary asks. A standard tells an institution what to aim for. The 5×5 Control Matrix™ measures whether it can demonstrate it has arrived.

Standards define what good governance should look like. We measure whether you can prove you have it.

Four differences that matter

- **Prescription versus evidence.** Standards prescribe practices and invite self-attestation against them. Our matrix reads what an institution has actually demonstrated, and grades the completeness of that demonstration. The unit is not “did you adopt the practice” but “can you prove it.”
- **Self-reported versus externally derived.** Conformance is typically self-declared. Our assessment is derived independently, from public disclosure, by a reviewer the institution did not engage — closer to a rating than to a compliance attestation.
- **The whole stack, not only the model.** Standards concentrate on model governance. Our matrix reads five ecosystems — power, compute, data centers, models, and agents —

surfacing where control thins at the infrastructure layers that governance frameworks rarely examine.

- **Capability versus auditable control.** The distinction at the heart of our method — a named capability is not the same as an auditable control mechanism — is one most standards have no language for. It is also the gap where institutional AI risk now lives.

We build on the standards — we do not replace them

An institution aligned to NIST or certified to ISO 42001 has done important work, and that work is a foundation, not a finish line. Policy on paper and control in practice are different achievements; the first is necessary for the second and does not guarantee it. We measure the second — across the full AI stack, and from the vantage point of an examiner rather than the institution itself.

A disclosure-grounded discipline

Our reading is deliberately grounded in what an institution discloses publicly — because disclosure is itself the thing under examination. What an institution will and will not say about its own AI control is evidence, and in our review of eighty leading institutions, most could say remarkably little. Absence of demonstrable control is not a verdict on any single firm; it is a measure of how far the industry still has to travel between governing AI in principle and commanding it in practice.

AI is a given. Control is not.™