AN EXECUTIVE BRIEFING

The AI Risk & Governance Review™

How executives defend their AI decisions when the board, the regulator, the acquirer, or the lawyer asks.

Stephen R. Jordan
Founder, SRJ Consulting & Services LLC

Volume III of The Operating Discipline for AI LibraryÅ

II

Having governance is not the same as being able to prove it.

Stephen R. Jordan · The AI Risk & Governance ReviewÅ

The question is arriving in five costumes.

Most leadership teams run disciplined businesses with real controls that have survived real audits. But somewhere ahead, on a date you do not control, someone with authority will ask you to prove your AI is governed. The only question is what you hand them.

01

The carrier
A D&O or cyber renewal questionnaire with a new AI section.

02

The customer
An enterprise vendor risk team gating the contract on an AI attestation.

03

The acquirer
A diligence list with a new section on AI governance — priced into the deal.

04

The regulator
An inquiry with a response window measured in days.

05

The board
A chair turning to the CEO and asking: who is accountable for this?

The average U.S. data breach now costs

$10.22M

Documented AI governance is on the right side of that ledger: AI/ML security insights save $223K per breach, governance technology $192K, documented AI policies $147K. **Ungoverned shadow AI adds $200K.**

IBM / Ponemon Institute, Cost of a Data Breach Report 2025

The Governance Gap

Standard governance is sound. It was simply built on assumptions AI quietly breaks.

Where the exposure quietly lives:

- Data leaves through doors your controls were never watching
- AI-assisted decisions blur who actually decided
- Vendors change models, terms, and sub-processors between your reviews
- A second rulebook arrived that never routes through your compliance calendar
- Your incident plan has no playbook for confidently wrong AI output

Governed on paper is not provable on request.

GOVERNED ON PAPER

- A generic AI policy nobody operates
- Vendor contracts reviewed once, at signing
- Collective accountability — “everyone owns AI”
- A narrative assembled after the question arrives

PROVABLE ON REQUEST

- A signed governance dossier per material use case
- A vendor risk inventory reviewed on a cadence
- One named executive accountable, in writing
- Documents that were waiting for the question

Additive, Not Replacement

II

AI governance does not replace the governance you have.

It covers what the governance you have was never designed to see.

Eight specific gaps. Each one invisible until someone with authority asks. Each one inexpensive to close on your own calendar. calendar.

Apply the discipline you already trust

to the one function that arrived without it.

01 **Named accountability**

A single Executive Sponsor. Not a committee. Not a shared mailbox. A name.

02 **Four risk categories**

Data, decisions, vendors, compliance. Every framework requirement lands in one.

03 **Tiered controls**

Oversight proportionate to impact — light where stakes are low, rigorous where they are not.

04 **Documentation that holds**

Contemporaneous records: dossiers, logs, minutes, sign-offs. Dated when it happened.

05 **A review cadence**

On the calendar. Recurring. Defensible to a board, an auditor, or a regulator.

Three businesses that thought they were covered

Composite cases from the book. Well-run companies. Strong controls. That is the point.

THE MANUFACTURER

ISO 9001. SOC 2 Type II. Active audit committee. Copilot and a production-planning AI passed every gate the company had — and none of the gates asked the AI questions. A customer questionnaire found nothing to send.

THE SERVICES FIRM

Fifteen clean D&O renewals. A candidate-evaluation AI, procured properly. Then a renewal questionnaire asked for the bias audit and the named accountable executive — while the market attached AI exclusions to the policy it counted on.

THE ACCOUNTING FIRM

Peer review on schedule. A compliance culture most firms envy. And staff using free-tier ChatGPT — sometimes with client information. No policy prohibited it, because no policy contemplated it.

Three signs your AI is ungoverned

01 **You could not produce the document.**

If the request arrived tomorrow — carrier, customer, acquirer, regulator — there is no dossier, no statement, no evidence package. There is a task force.

02 **No single name is accountable.**

Responsibility for AI is collective, informal, or unassigned. When something goes wrong at two in the afternoon, there is a debate about whose phone rings.

03 **Nobody knows which use cases are high-stakes.**

AI touching hiring, pricing, customers, or financial reporting carries the same oversight as AI drafting internal emails: none.

The Frameworks, Without the Fog

II

When the question arrives, it arrives wearing a framework's name.

You do not need a consultant's glossary to answer it.

ISO 42001. The NIST AI Risk Management Framework. The EU AI Act. Local Law 144. SR 11-7.
In plain English, sized for a business your size.

What the frameworks actually ask of you

ISO/IEC 42001

The management system standard — think ISO 9001, for AI.

It asks whether you run a system: leadership, risk assessment, assessment, controls, documentation, internal audit, management review, corrective action. Alignment first; first; certification is a commercial decision.

NIST AI RMF

Voluntary, American, practical. The language of customer questionnaires.

Four functions — govern, map, measure, manage — and seven trustworthiness characteristics: valid, safe, secure, accountable, explainable, privacy-enhanced, fair.

EU AI ACT

Law, not guidance — and it reaches U.S. companies serving EU serving EU customers.

Risk tiers with real dates: prohibited practices banned Feb 2025, GPAI obligations Aug 2025, the high-risk regime — hiring, lending, consequential decisions — from Aug 2026.

The Four AI Operational Risk Categories™

Every framework requirement, from every source, lands in one of four buckets you already understand.

01

Data

02

Decisions

03

Vendors

04

Compliance

Manage the four categories well, and you are most of the way to answering any framework anyone cites. The AI Governance Framework Crosswalk maps every requirement to the artifact that responds.

Data Risk

What leaves the building through AI tools, and under what protections.

Operating instrument: The AI Data Exposure Model

LEVEL 1 — Public	Already public or intended for publication. Straightforward use.
LEVEL 2 — Internal	Ordinary business information. Requires an enterprise contract: no training on prompts, defined retention, sub-processor disclosure.
LEVEL 3 — Confidential	Trade secrets included. Zero-retention configuration, access controls, exception logging.
LEVEL 4 — Regulated	PHI, GDPR data, cardholder data. Fully isolated environments with the regulatory paperwork.
LEVEL 5 — Restricted	Source code as the asset, deal workpapers, privilege. Never leaves infrastructure you control.

The rule employees memorize: nothing above Level 2 ever enters a free consumer AI tool.

Decision Risk

Who actually decided — and whether the oversight matches the stakes.

Operating instrument: The Decision Influence Matrix

LOW	User reviews their own output. Drafts, summaries, notes. Personal verification only.
MODERATE	Output goes beyond the user; reversible. A manager signs off, with evidence.
HIGH	Influences a consequential decision — hiring screens, pricing, financial reporting. Named reviewer, documented validation, audit trail.
CRITICAL	Legal or regulatory weight; limited recourse. Continuous human-in-the-loop, complete audit trail, bias audits where law requires.

Most businesses find two or three use cases at High or Critical with no controls at all. That is where the exposure lives.

CATEGORY 03

Vendor Dependency Risk

Your AI runs on someone else's model, under someone else's terms.

Operating instrument: The AI Vendor Risk Inventory

Training rights

Does the vendor train on your prompts? The single most consequential contract term.

Retention & residency

How long your data persists, and where it is processed and stored.

Sub-processors

Who else touches the data — with change-notice rights in in writing.

Caps & coverage

Indemnification and liability caps. What the vendor actually stands behind.

Framework position

ISO 42001, SOC 2, EU AI Act classification — and your resulting obligations.

Continuity owner

A named owner of the exit plan. Vendors get acquired, reprice, and sunset products.

Reviewed on a cadence — not just at signing — because AI vendor terms change between signings.

CATEGORY 04

Compliance & Financial Reporting Risk

The second rulebook — including the one your CFO personally certifies.

The map changes quarterly. An EU statute with extraterritorial reach. New York City's annual bias-audit law for hiring tools. Colorado, Texas, California, and Illinois legislating in waves. Insurance exclusions written specifically for generative AI. Agency enforcement from the FTC, EEOC, CFPB, and SEC.

THE ONE MOST TEAMS MISS

If AI touches anything that flows into your financial statements, your auditor is now interested — the AICPA published an AI checklist, the PCAOB put AI on its inspection focus, and under SOX 302 and 404 your certifying officers are personally attesting to controls that now have AI inside them.

You do not need to memorize the map. You need a discipline that tracks it — and a crosswalk that says which rules touch which use cases.

The 6-Step Review Process

First full Review: two weeks to two months. Every cycle after is faster.

01 Discovery

Confirm what AI is actually running, with named owners. owners. Shadow AI surfaces here.

02 Assessment

Score every use case with the Volume I and II instruments, so everything is comparable.

03 Mapping

Which framework requirements apply, and which artifacts respond. The fog clears here.

04 Risk Identification

A prioritized risk register per use case.

05 Target-State Design

What good looks like: controls, oversight, documentation, cadence.

06 Migration Planning

Named owners, due dates, dependencies — feeding the 90-day cycle.

Five dimensions. One number.

Oversight, accountability, controls, documentation, review cadence — each scored 1 to 5.

5	**Optimized**	The cycle compounds. Scores improve, findings decline, discipline extends automatically.
4	**Managed**	The cadence operates and produces evidence. An external party could be handed the dossiers today.
3	**Defined**	Named accountability, adopted controls, organized documentation. The installable target.
1–2	**Ad Hoc / Initial**	Where most companies honestly score on the first pass. A baseline, not a failure.

Boards respond well to a baseline with a credible plan. They respond badly to discovering there was never a baseline at all.

The dossier is what you hand them.

One document per material AI use case. Eight to fifteen pages. Ten sections, from named accountability through the Executive Sponsor's signature.

It is what the regulator opens in an inquiry. What the vendor risk team team reviews. What the diligence team reads. What your board sees. sees.

THE TEN SECTIONS

01 Identification & named accountability

02 Tools & vendors

03 Data flow & exposure level

04 Risk classification & tier

05 Framework alignment

06 Trustworthiness position

07 Human oversight & reviewers

08 Performance & friction

09 Risks & remediation

10 Sign-off

Could you answer these today, with documents?

01 Which use cases could influence a consequential decision about a person or a financial outcome — and what controls are active on each?

02 What company data is leaving through AI tools, and under what contractual protections?

03 If a regulator, customer, or carrier asked us to demonstrate our AI governance, what would we produce — and how fast?

04 How is this board exercising its oversight obligation for AI as a material risk category?

05 What AI incidents occurred in the trailing twelve months, and what corrective actions resulted?

Five for five with documents: you can skip the book. Two or three stomach-tightens: you now know which chapters to read first.

Three patterns decide how that room goes.

CONTEMPORANEOUS

beats reconstructed

A corrective action entry dated the day of the event carries weight. A narrative assembled by counsel six months later does not.

NAMED

beats collective

“The Executive Sponsor paused the use case at 2:47 PM on March 14th” is a strong sentence. “The team determined it was appropriate” is not.

DELIBERATE

beats reactive

A review that ran on a calendar, before anything went wrong, demonstrates governance. One that started the day after demonstrates damage control.

Every instrument in the book is built for a reader you hope never shows up. If that reader never arrives, the same records win renewals, contracts, and clean diligence.

What the work is worth

INSURANCE

ISO has published generative-AI exclusion endorsements; carriers attach them at renewal. A specialty market of affirmative AI coverage is emerging. coverage is emerging. Which side you land on is decided by the documentation this Review produces.

REVENUE

Enterprise customers gate contracts on AI attestations; the vendor questionnaire is the new front door. Companies that can produce the third-party the third-party governance statement close faster.

THE DEAL

Acquirers price what they find. Undocumented AI reads as unquantified liability — and unquantified liability comes out of the purchase price, becomes an escrow, or becomes a walked deal.

Governance proportionate to risk — not maximum governance. Sized to the business, paying for itself in premiums, in deals, and in the breach handled better or never had. handled better or never had.

Three blocks. Three completion checks.

Runs on named owners and existing meetings — not on consultants living in your conference room.

DAYS 1 – 30

Foundation

- Name the Executive Sponsor
- Approve the Accountability Matrix
- Update the policy + functional addendums
- Open the three governance logs
- First Steering Committee meeting

DAYS 31 – 60

Perimeter

- Run the 6-Step Review on priority use cases
- Draft the per-use-case dossiers
- Executive Sponsor sign-off
- Document the maturity score

DAYS 61 – 90

Scale

- Extend to second-tier use cases
- Launch the literacy program
- First management review
- Severity-1 tabletop exercise
- Deliver the board briefing

What this Review is not

A governance and risk discipline for operators — proportionate, not maximal.

Not a technology assessment

We don't evaluate which AI tools to buy. We make the AI you already run defensible.

Not maximum governance

Proportionate to your size and risk profile. The book says plainly what businesses your size can skip.

Not legal advice

It prepares the documents and discipline your counsel works from. It does not replace counsel.

Not a one-time project

An annual Review on a standing cadence — a discipline to operate, not a binder to shelve.

BUILT FOR

Who this Review is for

Written for executives, not engineers. There is no math in it.

ROLES

- Owners and presidents
- CFOs, COOs, and general counsels
- Boards and audit committees
- Operating partners overseeing AI risk

SECTORS

- Professional services and consulting
- Financial services, accounting, and legal
- Healthcare and insurance
- Manufacturing and distribution

Built for businesses of roughly twenty to one thousand employees — large enough that AI is inside the operation, and large enough that someone external will someone external will eventually ask hard questions about it.

ABOUT THE AUTHOR

Stephen R. Jordan

Founder, SRJ Consulting & Services LLC

Author of The AI Risk & Governance ReviewÁ , the third Volume in The Operating Discipline for AI LibraryÁ , a nine-volume canon on running AI as a permanent business function. Stephen also wrote The AI Business Enablement AuditÁ (Volume I) and The AI Readiness & Performance AssessmentÁ (Volume II), and is developing the remaining Volumes across the AI Business ServicesÁ and AI Risk Governance & SecurityÁ pillars.

Stephen draws on three decades of senior security and operations leadership across Citi, Intel, McAfee, and Optiv. His consulting work focuses on AI governance, operational discipline, and business performance for organizations from mid-market to large multinational conglomerates.

Based in Dallas–Fort Worth, working with executives nationwide.

BACKGROUND

Citi

Intel

McAfee

Optiv

When they ask,
what will you hand them?

Book a consultation with Stephen R. Jordan.

srjconsultingservices.com