CISO QUICK REVIEW EDITION

THE AI MANDATE

Why AI Governance Belongs in the CEO's Office, Not the IT Department

Manav Chadha

MBA · CISM · A|CISO · ISO 42001 Lead Auditor · ISO 27001 Lead Auditor

Launching June 12, 2026

What You Get	Reading Time
Why the CEO must own AI governance — not the CISO	2 min
The AI Data Leakage problem — sector by sector	3 min
The Semantic Leakage Prevention Prototype	2 min
The Three Red Lines — non-negotiable rules	2 min
Vendor Risk — questions and red flags	2 min
The Governance Cascade — why mandates die at Level 4	2 min
The 90-Day Sprint — CISO action items	2 min
Framework Alignment — ISO 42001, NIST, SOC 2	2 min
CISO Diagnostic — questions for your CEO	1 min

THE PROBLEM

Why This Book Exists

The most expensive sentence in corporate leadership.

85%	$4.88M	68%
of organizations have no formal AI policy ISACA 2025	average cost of a data breach IBM 2024	jump in third-party breach involvement Verizon 2024

The most expensive sentence in corporate leadership is not about a missed forecast or a failed product.

"The IT team handles all of that."

Every CISO reading this has heard it. Most have felt the frustration of it. You can see the risk clearly. The person with the authority to act believes the risk is managed because they have you.

This book was not written for you. It was written for the person above you. The one who needs to understand that AI governance is their accountability — not yours to manage alone and not theirs to delegate without consequence.

The author spent thirty years inside IT and ten years in cybersecurity. He built a semantic leakage prevention prototype. He has been in the server room when systems failed and in the boardroom when leaders did not have the answers they needed.

HOW TO USE THIS BOOK

For CISOs: This book is the most useful tool you have for the conversation you have been trying to have with your CEO. Hand it to them before you hand them another framework document. Stories move people. Frameworks do not.

The AI Governance Maturity Model

Level	State	What It Looks Like	Where Most Orgs Are
Level 1	Unaware	AI tools deployed. No policy. No owner. No audit trail.	Majority
Level 2	Reactive	Policy exists after an incident. Controls triggered by events.	Common
Level 3	Managed	Risk Register active. Vendors audited. Named owners.	Minority
Level 4	Proactive	Governance is a competitive advantage. Measurable and tested.	Rare

The gap between Level 1 and Level 3 is not a technology gap. It is a leadership decision. Organizations move from Level 1 to Level 3 when the CEO owns the mandate, not when the CISO writes better policies.

THE EMERGING RISK

AI Data Leakage

What your employees are doing right now without your knowledge.

Semantic leakage is what happens when confidential organizational data enters an AI model through ordinary employee use. Not through a breach. Not through a malicious actor. Through a professional trying to do their job faster.

Most employees do not know this is happening. They are not being negligent. They are using AI with context — without understanding what that context costs the organization.

The Data Being Fed Into Public AI Tools

Profession	What They Upload	Risk Tier
Legal	Case files, client communications, privileged documents	RESTRICTED
Finance	Revenue figures, projections, deal terms, personnel costs	RESTRICTED
HR	Personnel files, compensation data, performance records	RESTRICTED
Operations	Client names, project details, vendor contracts	CONFIDENTIAL
Executive	Board materials, M&A details, strategic plans, forecasts	RESTRICTED
IT / Security	Network diagrams, system context, vulnerability details	RESTRICTED

Regulatory Exposure

Regulation	AI Leakage Exposure	Potential Consequence
GDPR / PIPEDA	Personal data in AI prompts without consent	Regulatory fine, mandatory breach notification
HIPAA	Protected health information in public AI tools	Civil and criminal penalties
PCI DSS	Cardholder data in uncontrolled AI environment	Compliance failure, card brand penalties
SOX / CSOX	Financial data in unaudited AI pipeline	Internal control failure, CEO/CFO certification risk
Solicitor-Client	Legal strategy in public AI model	Privilege waiver, professional liability

FROM EXPERIENCE

I have a strong feeling that people are using AI for everything they should not. If something happens, the biggest question is who is responsible — because the source of truth could be missing, or who knows what data trained what model, by whom, and what it resulted in. — Manav Chadha

TECHNICAL CREDIBILITY

The Semantic Leakage Prevention Prototype

Not a framework. A working solution built to address a real problem.

Most AI governance frameworks were written by people who studied the problem. This prototype was built by someone who solved it. The prototype sits as a filter layer between the employee and the public AI tool.

The employee writes their prompt normally. Before it reaches the AI, the filter intercepts it, identifies sensitive entities — names, emails, account numbers, financial figures, client names — and replaces each with an anonymized placeholder token. The cleaned prompt goes to the AI. The output comes back through the filter, which maps the real values back. The employee receives a complete, usable response.

BEFORE — What the AI Receives	AFTER — What the AI Sees
Draft an email to Sarah Mitchell at Acme Corporation regarding the $2.4M contract renewal for Project Horizon. Her email is sarah@acmecorp.com. The delay was caused by Q3 revenue issues and Marcus Chen has resolved.	Draft an email to [CONTACT_1] at **[COMPANY_1] regarding the [VALUE_1]** renewal for [PROJECT_1]. Her email is [EMAIL_1]. The delay was caused by [PERIOD_1] issues and [CONTACT_2] has resolved.

✓ The employee gets what they need. The AI never sees the real data.

FOR CISOS

For CISOs: The prototype demonstrates that the technical gap is solvable. The governance gap — the absence of a CEO-level mandate requiring this kind of control — is the harder problem. Lead with the technical credibility when presenting to your board. Lead with the governance mandate when presenting to your CEO.

THE POLICY FOUNDATION

The Three Red Lines

Simple enough to memorize. Clear enough to enforce.

Rule	The Standard	Why It Works
Rule 1 Public Domain Standard	If you would not post this information publicly, you do not feed it to a public AI tool.	Gives employees a test they can apply in the moment. No training required. No ambiguity.
Rule 2 The Attachment Ban	No internal documents, client data, or proprietary content uploaded to any public AI tool.	Eliminates the highest-risk action in the AI workflow. The upload is where the most significant exposure occurs.
Rule 3 Zero Retention Verification	Before approval, obtain written contractual proof that the vendor does not retain prompts or data.	Not a privacy policy. A contract clause with consequences. Privacy policies change. Contract clauses have remedies.

THE CRITICAL DISTINCTION

A rule that comes from IT is optional. A rule that comes from the CEO is policy. That distinction is not subtle. It changes behavior at every level of the organization. The Three Red Lines must be issued by the CEO, with the CEO's signature, to every employee.

CISO ACTION ITEMS — RED LINES

- Draft the AI Acceptable Use Policy this week. One page. Three rules. CEO signature. Appendix C of the full book is a complete ready-to-use template.
- Build the approved AI tool inventory. Every tool, named Business Owner, named Technical Owner, data classification permitted, last review date.
- Send the Vendor AI Audit Questionnaire to your three most critical AI vendors. Appendix B of the full book is a complete 15-question template with a scoring guide.
- Confirm zero data retention clauses in writing for every AI tool handling internal or client data. Not the privacy policy. The contract.

THIRD-PARTY RISK

The Third-Party Trap

Your vendor's breach is your breach. The regulator does not distinguish.

Third-party involvement in data breaches increased by 68 percent in a single year according to Verizon's 2024 DBIR, which analyzed more than 10,000 confirmed incidents. The governance gap enabling it is consistent: organizations collect vendor security documentation and file it without reading it.

The Vendor Assessment Standard

Question	Acceptable Answer	Red Flag
SOC 2 Type 2 report	Current report covering last 12 months, available within 5 days	Type 1 only, working on it, or unavailable
Data retention policy	Written clause: zero retention of prompts, contractually enforceable	Verbal assurance or privacy policy reference only
Breach notification	24-hour commitment in contract, with consequences	Reasonable time, as soon as practicable, not in writing
Penetration testing	Annual, executive summary available on request	Not conducted, no schedule, results unavailable
Sub-processors	Full list disclosed, data residency confirmed	Undisclosed, or subject to change without notice

The Contract Clauses That Cannot Be Negotiated

⚠ 24-hour breach notification — under most regulations you have 72 hours to report. Your vendor must notify you within 24 hours to give you adequate response time.

⚠ Right to Audit — the right to conduct or commission an independent security assessment. Vendors who refuse this clause have something an audit would find.

⚠ Zero data retention — prompts and data submitted are not retained, not used for model training, deleted on request. In writing. In the contract.

⚠ Data residency — explicit geographic commitment for where your data is stored. Not subject to change without your consent.

THE HALO EFFECT

Brand is not a security control. A vendor whose 32-bit architecture cannot support a modern enterprise migration will tell you it supports everything — until you ask the specific technical

> *questions that reveal the gap. The quality of a vendor's sales presentation tells you exactly nothing about the quality of their security controls.*

ENTERPRISE GOVERNANCE

The Governance Cascade

Why CEO mandates die at Level 4 — and how to fix it.

Level	Role	What They Must Own	Common Failure Mode
Level 1	CEO	Non-negotiable standard. Signed mandate. Quarterly review chair.	Sets mandate but treats it as a one-time act. Does not follow up.
Level 2	CIO / CTO	Architecture decisions. Approved tool inventory. Technical controls.	Translates mandate into technology without making it operational for business units.
Level 3	CISO / CRO	Control validation. Vendor audit. Unfiltered upward reporting.	Softens bad news before it reaches Level 1.
Level 4	Operations / Divisions	Policy enforcement. Named AI tool owners. Compliance reporting.	Deprioritizes governance when it competes with quarterly targets. MOST COMMON FAILURE POINT.
Level 5	Service Teams / Frontline	Daily policy compliance. Escalation when uncertain.	No named escalation path. No training. Default to individual judgment.

THE CASCADE PRINCIPLE

The gap between the CEO's mandate and the frontline employee's behavior is not a technology gap. It is a translation gap. Level 4 receives a policy document. What they need is a named tool list, a named owner structure, a named escalation path, and a named consequence for non-compliance.

The Test That Reveals the Gap

Ask your CEO this question in your next one-on-one. Do not prepare them. That is the point.

"Without calling anyone on the IT team — right now, in this conversation — can you name which AI tools your operations teams are using that were not formally approved?"

If they cannot answer, or do not know who to call, the cascade has broken somewhere between Level 1 and Level 4. The 90-Day Sprint below is how to fix it.

THE ACTION PLAN

The 90-Day Governance Sprint

From theoretical to operational in ninety days.

Phase	Days	CISO-Led Actions	CEO-Required Actions
Phase 1 Foundation	1–30	Build AI tool inventory. Assign named owners. Assess BCP. Draft AI mandate.	Sign and issue mandate to all staff. Authorize Risk Register review. Run tabletop.
Phase 2 Vendor & Data	31–60	Send Vendor Audit Questionnaire. Review SOC 2 reports. Verify retention clauses. Draft Data Classification Policy.	Approve Data Classification Policy. Attend one training session. Review vendor risk summary.
Phase 3 Policy & Monitor	61–90	Finalize AI Acceptable Use Policy. Stand up monitoring function. Run AI tabletop. Prepare first Governance Review.	Chair first Governance Review. Sign off on minutes. Lock quarterly governance calendar.

CISO — WHAT TO PRESENT TO YOUR CEO AFTER THEY READ THIS BOOK

- One page. Three Red Lines. Your signature request on the AI mandate. Do this in the first week.
- The AI tool inventory — what is approved, what is in use without approval, who owns each one. Present this as a risk briefing, not a technical report.
- The vendor risk summary — which critical vendors have current SOC 2 Type 2 reports, which do not, and what you need the CEO to authorize.
- The 90-Day Sprint calendar — dates, owners, decision points that require CEO involvement. Make it easy to say yes.
- The first Governance Review agenda — ninety minutes with the full leadership team. Frame it as a business risk meeting, not an IT status update.

FOR YOUR CEO CONVERSATION

The CISO Diagnostic

Questions that reveal governance posture without a single dashboard.

These questions are designed to be asked of your CEO — or used by your CEO to self-assess — without preparation. Preparation defeats the purpose. The honest answer is the useful one.

⚠ Can someone produce the Risk Register in five minutes right now?

⚠ Which AI tools are your operations teams using that were not formally approved?

⚠ When did you last read — not receive — the most recent vendor security report?

⚠ Who is the named owner of each AI tool your organization uses?

⚠ What happens to your operations if your primary AI tool is unavailable for 48 hours?

⚠ Has your BCP been tested in a live exercise in the last twelve months?

⚠ What percentage of your employees have been trained on AI data handling this year?

HOW TO USE THIS

Send your CEO this review edition before your next governance conversation. Ask them to answer the seven questions above before they read further. Their answers — honest, unfiltered — will be the most accurate governance assessment you have conducted. No consultant required.

COMPLIANCE ALIGNMENT

Framework Alignment Reference

How The AI Mandate maps to the standards your organization already uses.

Book Concept	ISO 42001	ISO 27001	NIST AI RMF	SOC 2
CEO Mandate & Accountability	4.1, 5.1	5.1 Leadership	GOVERN 1.0	CC1.1
AI Tool Inventory & Ownership	8.1 Operational	A.8.1 Assets	MAP 1.0, 2.0	CC6.1
Three Red Lines / Data Policy	8.4 AI Data	A.8.2 Classification	MAP 5.0	CC6.1, CC6.7
Vendor AI Audit	8.5 Third-Party	A.15 Suppliers	MAP 3.0	CC9.1, CC9.2
Governance Cascade	5.3 Roles	6.1 Risk Treatment	GOVERN 4.0	CC1.3
Ethics / Transparency	6.1, 8.6 Ethics	A.6.1 Roles	GOVERN 6.0	CC2.1
Incident Response	8.9 Incidents	A.16 Incidents	RESPOND 1.0	CC7.3, CC7.4
90-Day Sprint / Monitoring	9.1, 10.1	9.1 Monitoring	MANAGE 4.0	CC4.1

THE CERTIFICATION OPPORTUNITY

Organizations that implement the 90-Day Sprint and maintain the quarterly governance calendar are building the operational evidence base required for ISO 42001 certification. The six templates in the full book appendix are designed to produce the documented evidence that an ISO 42001 audit requires.

THE FULL BOOK INCLUDES SIX PROFESSIONAL TEMPLATES

- Appendix A — IT Risk Assessment Template (complete, fillable)
- Appendix B — Vendor AI Audit Questionnaire (15 questions, scoring guide, decision framework)
- Appendix C — AI Acceptable Use Policy (ready to adapt, CEO sign-off page included)
- Appendix D — Data Classification Policy (four-tier framework, plain language)
- Appendix E — Quarterly Governance Review Agenda (90-minute structured format, escalation triggers)
- Appendix F — Incident Response Checklist (72-hour, step-by-step, named owners at every step)

Share This With Your CEO

If what you have read here reflects the governance gaps you have been trying to close, the full book gives you the language, the stories, and the tools to make the case — in a format your CEO will actually read.

Twenty-eight minutes. A 90-day action plan. Six professional templates they can hand back to you and say: build this.

Most AI governance books were written for IT. This one was written for the person IT reports to.

> **THE AI MANDATE**
>
> *Available June 12, 2026. For every leader willing to be accountable before the crisis demands it. — Manav Chadha, The AI Mandate*

About the Author

Manav Chadha is a technology leader, cybersecurity practitioner, and governance advisor with thirty years of IT experience and ten years of specialized cybersecurity practice. He holds an MBA, CISM, A|CISO, ISO 42001 Lead Auditor, ISO 27001 Lead Auditor, AI Governance Professional, and SOC 2 Auditor credentials. He is the founder of TrustGate Security Inc. and is based in Edmonton, Alberta, Canada.