THE NEW BOOK | LAUNCHING 12 JUNE 2026

THE AI MANDATE

Why AI governance belongs in the CEO's office, not the IT department.

MANAV CHADHA

MBA • CISM • ASSOCIATE C|CISO • CCEP

Copyright and Disclaimer

The semantic leakage prevention framework and prototype described in this book are the original work of Manav Chadha, developed independently prior to publication.

Every statistic referenced in this book is drawn from publicly available research and is attributed to its source. The author has no commercial or professional affiliation with any organization cited.

Every story in this book comes from real professional experience. Where details could identify a specific organization or individual, they have been changed to protect confidentiality. No story is invented. Every lesson is real.

This book reflects the author's experience and professional judgment accumulated over 25 years in IT and over a decade in cybersecurity. It is not legal, regulatory, financial, or professional advice. For decisions carrying significant organizational consequence, work with a qualified professional who understands your specific situation.

Dedication

For my mother.

She never attended a governance conference. She never read a compliance framework. She did not know what a Risk Register was. But she understood accountability, honesty, and doing the right thing better than anyone I have ever met in 25 years of professional life.

I was finishing this book when I lost her in 2025. Somewhere in those early weeks of grief, I realized that every principle in this book, the accountability, the transparency, the insistence on doing the right thing before the crisis demands it, she had already lived. Long before I had words for any of it.

She was the best teacher I ever had. She taught me by doing, not by instructing. By choosing, not by lecturing. By being honest when a comfortable lie would have been so much easier.

This book carries her values on every page, even when it does not say her name.

For my wife and children.

You gave me the space, the time, the belief, and the patience that a book requires. Every morning I spent writing, you carried more than your share so I could carry this. I do not take that lightly. I never will.

For every leader willing to be accountable before the crisis demands it.

This book was written for you. Not for the leader who already has everything right. For the one who knows something is missing and is ready to do something about it.

Acknowledgments

To the senior leaders and mentors who challenged my thinking and trusted me enough to push back when they disagreed. You made every argument in this book sharper.

To the teammates who were in the server room with me through the long nights, manually working through infrastructure that had been neglected for years. That experience became the first chapter of this book.

To my advance readers, whose honest feedback, especially the critical feedback, made this book better than it would have been.

To my family, who gave me the most valuable resource any author can have: time, patience, and the belief that what I was building mattered before there was any evidence it would.

And to every professional I have mentored over the years, more than two hundred of you now. You taught me as much as I taught you.

FOREWORD

A Word Before You Begin

I have known Manav for almost ten years now.

We first crossed paths at BSides Edmonton. If you have never been to a BSides conference, it is worth explaining what makes it different. It is not a corporate event. There are no polished vendor booths and no keynotes designed to sell you something. It is a grassroots gathering of people who work in security because they genuinely care about it, people who show up on a weekend to talk honestly about what the work actually looks like. The conversations at BSides are different from the ones at bigger conferences. People talk about what went wrong, not just what went right. They talk about the calls they had to make without enough information and the situations they walked into that no training had prepared them for.

Manav fits that environment. He was not there to network or to be seen. He was there because he was thinking hard about a problem and wanted to talk to people who were thinking hard about the same thing. Manav is a go getter and has always been motivated to learn, engage and meaningfully contribute to the field of Information Technology and Cyber Security.

Over the years that followed we kept finding ourselves in the same rooms, LinkedIn, local IT events around Edmonton, cyber security conferences and meetups where the real conversations happen in the hallway between sessions, not in the auditorium during them. The kind of professional relationship that builds slowly and honestly until you realize you genuinely trust what the other person sees.

When he told me he was writing a book on AI governance for senior leaders, I was not surprised. I had watched him wrestle with this problem for years. Not as an abstract exercise but as someone who had been inside organizations where the gap between leadership intent and operational reality was producing real and measurable damage. He talked about it the way practitioners talk about things that genuinely bother them, with a specific kind of frustration that only comes from having tried to fix something and kept hitting the same wall.

That frustration is what makes this book worth reading.

I have worked across a lot of sectors in my career. Capital Power, where critical infrastructure security competed for attention against power generation performance and market pressures. Alberta Health Services, where patient data governance had to find its place among the daily operational demands of one of the country's largest health care organizations. Enbridge, where the complexity and scale of a major energy company means that security decisions made at the operational level carry consequences that can take months to surface at the leadership level. And years of teaching at Northern Alberta Institute of Technology and Concordia University of Edmonton, where I have watched students learn the technical side of this work and

then struggle to find the words that make a boardroom care about it.

In every one of those environments I have seen the same thing. Technical people know where the gaps are. They can see the risks clearly. What they often cannot do is get the right person to pay attention before something goes wrong. That right person, the CEO, the board member, the executive with the actual authority to change behavior across the whole organization, usually operates on a comfortable assumption: that because they have a cyber security team, the cyber security team has it handled.

Sometimes they do. But in the age of AI, the risk is no longer sitting only in your infrastructure. It is sitting in every employee with a laptop and access to a public AI tool who is trying to do their job faster and has no idea what happens to the confidential data they are typing into that tool. That is not an IT problem. It is a leadership problem, and most organizations have not yet recognized that it belongs on the CEO's desk rather than the CISO's.

Manav has recognized it. And most importantly, he has found a way to explain it that does not require the person reading it to have spent fifteen years in security to understand.

I want to say something honest about what it is like to be a security practitioner trying to make this case to senior leadership. It is genuinely difficult. Not because executives are indifferent. Most of the ones I have worked with are sharp, engaged, and serious about their responsibilities. The difficulty is translation. Our profession has spent decades speaking in a language that excludes the very people with the authority to act. We talk about threat vectors, attack surfaces and maturity frameworks, but somewhere in that language the urgency disappears. The person across the table hears the word governance and thinks paperwork. They hear risk register and think quarterly compliance box-ticking. They do not hear what it costs when it breaks.

This book speaks in a different language. It uses stories. Real ones, from someone who was actually in the room when things went wrong. The firmware story in Chapter 1 will be immediately recognizable to anyone who has managed infrastructure through an unplanned outage. The vendor rejection story, the dashboard that was telling a different story than reality, the semantic leakage prototype built because no one else was building it, these are not hypothetical scenarios constructed to make a point. They are the kind of things you talk about with a colleague you trust, late in a conference day, when the formal presentations are done.

The prototype is the detail that sets this book apart from everything else being written on AI governance right now. Most frameworks in this space were produced by people who studied the problem. Manav built something to address it. That distinction matters, and readers with a technical background will feel it from the first chapter.

The most important conversation in security is not the technical one. It is the conversation with the person who has the authority to act and who currently believes, incorrectly, that someone else is handling the risk. I have been trying to have that conversation in boardrooms and classrooms for years. Most of the time, the tools available to me have been frameworks and compliance requirements and documented standards, useful but not

always persuasive.

This book is a better tool. It gives practitioners language and stories they can put in front of their leadership. It gives leaders the questions they should be asking before those questions get asked by a regulator or a board or a headline. And it does both of those things without talking down to either audience.

I have known Manav for almost ten years. I have watched him think about these problems carefully and honestly across a lot of conversations in a lot of different settings. The AI Mandate is the book that thinking deserved to become.

Read it. Give it to your CEO. And if you are the CEO, read it yourself, starting with Chapter 1, before you hand it to anyone else.

Prashant

Senior Cyber Security Advisor, Enbridge
MSc • CISSP • CCSP • Associate CISO • GPEN • GICSP
Edmonton, Alberta • June 2026
linkedin.com/in/prashantprofile

Contents

PREFACE

The Delegation Trap

Why this book exists, who it is for, and the sentence that costs more than any missed forecast.

I did not plan to write this book.

I planned to keep doing what I had been doing for over twenty years, walking into organizations, finding the gaps between what leadership believed was happening and what was actually happening, and helping close them. Quietly. Practically. Without a book attached to it.

But somewhere in the middle of my career, and with increasing urgency in the last ten years, I kept running into the same wall. Not a technical wall. A human one.

The technical problems were always solvable. The governance gaps were always findable. What was not always possible was getting the right person in the right seat to own the accountability for fixing them.

"The IT team handles all of that."

I have heard that sentence from CEOs running organizations with thousands of employees. From board members who sat on audit committees. From senior leaders who had spent decades building genuinely impressive organizations and who genuinely believed that because they had delegated the technology, the technology risk had been delegated with it.

It had not. The accountability never moved. It sat on their desk the entire time, waiting quietly for the moment the crisis would make it visible.

This book is what I wish I could have handed every one of those leaders before that moment arrived.

Why I Am the Person Writing This

I want to be honest from the start. I am not the most prominent voice in AI governance. What I have is 25 years of being inside technology organizations when things went wrong. Not reading about what went wrong. Being in the room.

I have led enterprise governance programs that passed international audits. I have been inside organizations when their systems failed catastrophically. I have built a semantic leakage prevention prototype because I watched confidential data entering public AI tools every day without any governance behind it. I have coached more than two hundred professionals in IT, security, and governance.

And I have made mistakes. Some of them are in this book. Not as confessions but as lessons.

Why June 12th

I chose the publication date before I wrote the first word. June 12, 2026 is the one-year anniversary of losing my mother.

She never worked in technology. She did not know what a governance framework was. But she understood accountability better than anyone I have ever met in a boardroom. She built businesses on honesty. She made mistakes and her response to every mistake was to find the root cause, learn the lesson so deeply that it changed her behavior, and make sure the people around her learned it too.

I did not realize until I was well into writing this book that the values I was trying to articulate, accountability, transparency, doing the right thing before the crisis forces it, were things she had simply lived. Long before I had words for them.

This book is dedicated to her. To my wife and children who gave me the space to write it. And to every leader who wants to do the right thing and just needs someone to tell them where to start.

THE PROMISE OF THIS BOOK

You can delegate the management of risk. You cannot delegate the accountability for its consequences. When something goes wrong, the board calls the CEO. Not the CTO. Not the CISO. You.

THE AI MANDATE

A Statement of Executive Accountability

The AI Mandate is a commitment made by the senior leader of an organization to own, govern, and be accountable for how artificial intelligence is used within it.

By leading this organization I accept the following:

01

AI governance is my accountability.
I cannot delegate it completely. I can delegate the management. I retain the accountability for the outcome.

02

Every AI tool in active use will have a named owner.
Not a department. A person. If something goes wrong, there is one name attached to every AI tool in this organization.

03

No AI tool handling our data will be used without written proof of zero retention.
Not a privacy policy. A contract clause with consequences. Before any AI tool is approved, I will have written proof that our data is not retained.

04

My employees will know the rules because I will tell them, not IT.
A mandate from the CEO is policy. A mandate from IT is a suggestion.

05

I will chair a formal AI Governance Review every quarter.
I will read it. I will act on what it tells me. I will treat it as a leadership responsibility, not a reporting exercise.

This is The AI Mandate.

Signature	Date
Full Name	Title
Organization	

CHAPTER 1

The Leadership Mandate

You can delegate the task. You can never delegate the liability.

Here is the thing nobody tells you when you accept a senior leadership role.

They tell you about the strategy, the vision, the revenue targets. Nobody tells you that every single one of those things runs on technology. And that when the technology fails, nobody asks the CTO first.

They ask you.

Five Days That Changed Everything

We had a power outage. Should have been nothing. Power goes down, backup systems hold, power comes back, everyone moves on. That is how it works when an organization actually maintains its infrastructure.

We had not maintained ours. Not for years. When the backbone switches tried to restart, they failed. The firmware had not been updated in years. No documentation of when it was last done. No named owner. No process. We had to contact technical support, manually load firmware onto USB drives, and flash every switch one by one. Five days. No users could log in. No work could be done. No revenue being generated.

The CEO walked in. Measured. But the questions he asked made everything clear.

What happened?

What is our Recovery Point Objective, how much data did we lose?

What is our Recovery Time Objective, when can I tell the business we are back?

He was not asking about technology. He was asking about accountability. And the answers we gave him, incomplete, uncertain, still developing, told him everything he needed to know about how this organization had been governing its technology.

FROM THE TRENCHES

Five days of business downtime. The engineer who skipped the maintenance is long gone. The CEO still had to explain it to the board. That is what accountability without governance looks like up close.

The MDM Project: When Accountability Gets Lost in Translation

We were doing a Mobile Device Management project. The organization had brought in an external contractor to lead it. What happened next is something I have seen in some form in almost every large project I have been

part of.

The contractor believed one thing. The project manager believed another. The technical architect had a third interpretation. And the leader who had approved the project and considered it handled was sitting above all of them, unaware that none of the three versions of reality matched each other.

There was no communication between them. No collaboration. No single source of truth. So when the deployment team needed guidance, there was no guidance to give. No documentation. No direction. Just confusion at every level and a project that ground to a halt.

The look on the leader's face when this was surfaced was one I have seen many times since. Not anger, exactly. Something closer to the realization that the gap between what they had believed was happening and what was actually happening was their responsibility to close. One person lost their job. The rest of us learned something that no governance framework teaches: the gap between a leader's intention and an organization's execution is exactly the size of the accountability structures that fill it.

FROM EXPERIENCE

The gap between what a leader believes is happening and what is actually happening is always filled with something. Either it is filled with clear accountability structures and honest upward reporting. Or it is filled with confusion, assumption, and deferred problems.

What I Ask for on Day One

When I walk into an organization to assess its governance posture, I do not read the policy document they hand me at reception. I want evidence. Show me the Risk Register. Not a presentation about risk management. The actual register, with timestamps, with names on every line item.

If any request is met with hesitation, if anyone needs time to compile the information, I already know what I need to know. The organization is operating on hope rather than governance.

Hope is the worst disaster recovery strategy in existence.

Three Questions Every CEO Should Ask Today

Question	What the Answer Reveals
Can someone produce the Risk Register right now, without preparation?	Whether governance is operational or theoretical.
Which AI tools are employees using today that were not formally approved?	Whether your AI governance is real. If nobody knows, the tools are being used without rules.

Question	What the Answer Reveals
If primary systems went offline in the next hour, who are the first three people you call?	Whether your crisis response is a plan or just an intention.

MANAV CHADHA

Proper delegation is not walking away. It is delegating the task and staying accountable for the outcome. That is a leadership decision you make before the crisis. Not during it.

WHAT TO DO NEXT

Run the three questions in your next leadership meeting. Watch who hesitates. That hesitation is your gap.

Assign a named owner to every AI tool in use this week. Not a department. A person with a name.

Get the Risk Register in front of you before the next quarter ends. If it takes more than five minutes to produce, it needs rebuilding.

Send one email to your technology leadership this week: I want to understand our AI governance posture. Schedule thirty minutes with me.

CHAPTER 2

The ROI of Trust

Governance is not your biggest cost. It is your most powerful competitive advantage.

Most governance conversations start in the wrong place. They start with risk. That is the right conversation to have, but it is not the only one. And for many leaders, it is not the most persuasive one.

The CMMI Story: When Documentation Becomes Competitive Advantage

We were working on a CMMI Level 5 assessment. The Capability Maturity Model Integration measures an organization's process maturity. Level 5 is the highest level. Getting there means proving not just that your processes exist, but that they are consistently followed, continuously improved, and evidenced at every level of the organization.

When I came into the project, the scores were poor. Not because the organization lacked capability. Because nobody had organized the evidence. The knowledge was in people's heads, in email threads, in shared drives nobody had properly indexed. The work was being done. The evidence of it was nowhere to be found.

I started going through every team, every process, every control. Asking them to show me not just what they did but how they could prove it. The process was slow. Sometimes frustrating. There were moments where I found entire control areas where the evidence should have been and was simply not there.

But we got there. We brought the scores up. We presented the evidence. We passed. And what happened afterward was something I have seen replicated in every governance improvement program since. The satisfaction was not primarily external. It was internal. Our site had performed significantly better than others. We now had a documented baseline. A starting point. If something happened tomorrow, we knew exactly where we were, what we had, and how to recover.

FROM THE CMMI EXPERIENCE

The best time to build your governance baseline is before you need it. The second best time is now. A documented starting point is worth more than any number of undocumented good intentions.

The Vendor I Said No To

I was evaluating a technology vendor for a client organization. The product was genuinely impressive. The sales team was polished. Every person in the room was ready to sign.

I asked one question: can you produce your SOC 2 Type 2 report? A SOC 2 Type 2 report is not a policy document or a marketing claim. It is evidence that an organization's security controls have been operating consistently over a defined period, verified by an independent auditor.

They had a Type 1 report only. A snapshot of design, not proof of consistent operation. When I pressed, they said they were working on a Type 2. When I asked about their data retention policy for client prompts, they said they would get back to me.

We did not approve the vendor.

Some time later, that vendor experienced a data exposure incident that affected a number of their clients. My client was protected. The governance work had protected them.

MANAV CHADHA

A vendor who hesitates to answer how they protect your data, that hesitation is their answer. Brand is not a security control. A professional sales presentation tells you exactly nothing about the quality of security controls.

WHAT TO DO NEXT

Ask your procurement team: in our last five contract negotiations, were we ever asked to produce security certification evidence? What did we provide?

Request and actually read, not file, the most recent SOC 2 Type 2 report from every critical vendor.

Treat every passed audit and approved certification as a commercial asset. Put it in client proposals. Governance that nobody knows about cannot win contracts.

CHAPTER 3

Defining Your Red Lines

The most dangerous risks are the ones that look like convenience.

Let me tell you about something I believe is happening in your organization right now. Today. Without your knowledge or your permission.

Employees are feeding confidential data into public AI tools. Not because they are careless. Not because they are trying to cause harm. Because AI tools give better answers when they have more context, and nobody told them where the line was.

I have not personally witnessed someone deliberately decide to expose restricted data. What I have observed, repeatedly, is employees who are highly dependent on AI tools putting far more information into them than is safe, because more information produces better outputs and nobody told them where the line was. People are using AI for everything they should not. I do not call this responsible AI. Responsible AI means knowing what you are doing, understanding the implications, and making deliberate choices about what goes into these systems.

And then the biggest question becomes: if something happens downstream as a result of that data exposure, who is responsible? The source of truth could be missing. Who knows what data trained what model, by whom, and what it resulted in? That is dangerous.

What goes in may stay in.

The Three Rules That Cannot Be Negotiated

Rule One: The Public Domain Standard.

If you would not post this information on your public LinkedIn profile, you do not feed it to a public AI tool. No exceptions. This rule prevents the majority of data exposure incidents because it gives employees a concrete test they can apply in the moment.

Rule Two: The Attachment Ban.

No internal documents, financial data, client information, or proprietary content gets uploaded to a public AI tool. The moment you upload a file, you surrender custody of its contents.

Rule Three: The Zero Retention Verification.

Before any AI tool is approved for organizational use, obtain written contractual proof that the vendor does not retain your prompts or data. Not a privacy policy. A contract clause with consequences.

MANAV CHADHA

What you feed into AI is your responsibility. Not the vendor's. The source of truth for AI governance is not the vendor's terms of service. It is the policy your CEO signs.

WHAT TO DO NEXT

Send one email to your entire organization today: our AI data handling policy is rule one, rule two, rule three. It comes from me. It starts today. That email is governance.

Survey every department head this week: what AI tools is your team using? Build the inventory before you build the policy.

Before your next vendor renewal, ask in writing: do you retain our prompts or data? File the written answer with the contract.

INSIDE THE PROTOTYPE

How Semantic Leakage Prevention Works in Practice

Let me be specific about what I built. Not theoretical. Not aspirational. Here is the actual mechanism.

First, a word on the term. I use semantic leakage to describe a specific organizational problem: meaningful, sensitive information leaving your organization through the ordinary use of public AI tools, one prompt at a time. Researchers sometimes use the same phrase to describe a different phenomenon inside AI models themselves. In this book, the term means exactly one thing: your data, carrying your meaning, ending up where you no longer control it.

The prototype works as a filter layer between your employee and the AI tool. The employee writes their prompt. Before it reaches the AI, the filter intercepts it. It scans for sensitive entities, names, email addresses, account numbers, financial figures, client names, and replaces each with an anonymized placeholder token. The cleaned version goes to the AI. The output comes back. The real values are mapped back. The employee receives a complete, usable response.

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

Before: What the Employee Types	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.	Draft an email to [CONTACT_1] at [COMPANY_1] regarding the [VALUE_1] renewal for [PROJECT_1]. Her email is [EMAIL_1].

What This Is and What It Is Not

I want to be honest about where this sits in the landscape, because a book about honest governance should model it. Tools in this category exist. Data loss prevention platforms have intercepted sensitive content for years, and AI gateway products are emerging that filter prompts at the network level. I did not invent the idea of stopping sensitive data at a boundary.

What I built reflects a specific conviction: that the interception has to happen before the data leaves the employee's hands, that the employee should not have to change how they work, and that the response they get back must be complete and usable, or they will simply route around the control. Any control that makes the work harder gets bypassed. I have watched that happen with every inconvenient security tool of the last two decades.

And it has limits. Pattern recognition is not perfect. A determined employee can describe sensitive context in ways no filter will catch. A document summarized in the employee's own words still carries meaning no tokenizer can mask. No technical layer, mine included, replaces the policy, the training, and the leadership mandate that the rest of this book is about. The prototype closes the accidental path. Only governance closes the deliberate one.

WHAT I LEARNED

The technical solution needs the leadership mandate. You cannot code your way out of a behavior problem. The prototype stops the data from leaving. The policy stops the behavior that would send it. Both are required. Neither is sufficient alone.

THE NEW TERRAIN

What Makes AI Risk Different

Why the governance you already have is necessary, but no longer sufficient.

Much of what I have described so far is good technology governance. Named owners. Documented systems. Honest reporting. Tested recovery. These principles are not new, and a fair reader might ask what makes this a book about AI specifically rather than a book about technology leadership.

Here is the answer. The principles are the same. The terrain is not. AI introduces categories of risk that traditional technology governance was never designed to handle, and a leader who applies only the old playbook will miss them. These are the ones that belong on your agenda now.

Generative AI and the Confidentiality Problem

In my experience, the distinction is this. Traditional software does what you tell it. Generative AI produces new content based on patterns, and it does so using whatever you feed it. The risk, as I see it, is no longer only that a system gets breached. It is that your own employees may hand over confidential information voluntarily, in the ordinary course of trying to work faster, to a system that could retain and learn from it. This is the problem the prototype was built to address, and in my view it is the most immediate AI risk most organizations face today.

Agentic AI and the Accountability Problem

The newer AI systems, as I understand the direction of travel, do not just answer. They act. They can book, send, purchase, schedule, and execute across connected systems with limited human involvement. This is often called agentic AI, and to my mind it raises a governance question that did not exist before: when an autonomous system takes an action that causes harm, who is accountable? The vendor who built it? The employee who deployed it? The leader who approved its use? My own view is that the accountability sits exactly where this book has said it sits all along, with the leader, which is precisely why I believe agentic AI should not be adopted without a governance decision made at the top.

Model Governance and the Black Box Problem

With traditional software, you can usually trace a decision to a rule someone wrote. With AI models, my understanding is that you often cannot. The model produces an output and the reasoning is not fully inspectable. For a leader, this suggests you must govern not only the data going in but the decisions coming out. Who validates that the model's outputs are accurate, fair, and appropriate? Who is accountable when a model makes a confident recommendation that turns out to be wrong? Model governance, as I think of it, is the

discipline of answering those questions before the model is trusted with a decision that matters.

AI Regulation and the Compliance Horizon

AI regulation appears to be arriving. Initiatives such as the European Union's AI Act, emerging frameworks in North America, and the international standard ISO 42001 for AI management systems are, as I read them, early signals of a regulatory environment that will only grow. I would encourage any leader to verify the current status and specifics of these with their own legal and compliance advisors, as this area is moving quickly. My broader point is one of pattern rather than detail: the organizations that build governance now tend to adapt to regulation as a formality, and the organizations that wait tend to scramble to retrofit it under deadline pressure. I have watched the second pattern play out with every major compliance wave of the last two decades, and it has nearly always been more expensive than building it early.

AI Vendor Management and the Supply Chain Problem

In my experience, most organizations do not build their own AI. They buy it, often embedded inside tools they already use. This suggests that your AI risk is largely inherited from your vendors, and many organizations I have seen have little visibility into which of their existing tools have quietly added AI features that now process their data differently than they did a year ago. The Vendor AI Audit Questionnaire in Appendix B exists for exactly this reason. The question, as I see it, is no longer only what does this vendor do with our data. It is now also what does this vendor's AI do with our data, and has that changed since we signed.

THE CORE OF IT

The governance principles in this book are not new. The reason they matter more now is that AI has raised the stakes and shortened the timeline. The same accountability that has always belonged to the leader now governs systems that act on their own, learn from what they are given, and make decisions no one can fully inspect. The mandate has not changed. The cost of ignoring it has.

CHAPTER 4

The Risk Radar

A metric with no source of truth is not a metric. It is a guess dressed up as a number.

I want to tell you about a walk I took with a senior leader. Not a strategy walk. A literal walk through a building, from the boardroom down to the people who actually managed the technology every day. His dashboard said everything was fine. He had been making decisions based on that dashboard for months. It did not reflect reality.

The Story Behind the Green Lights

When I took him to the teams who were actually running the systems, what they described had nothing to do with the dashboard. Systems not updated in years. Processes that existed on paper but not in practice. Technical debt accumulating quietly, deprioritized quarter after quarter.

He was quiet for most of the walk. Processing. And at the end, he said something I have used in conversations with senior leaders ever since.

"I have been making decisions based on a story that was not true."

The Ego That Delayed Everything

Leadership changed in an organization I was working with after a merger and acquisition. The new leader had their own perspective on what information they needed before making decisions.

I had done the work. The numbers I was presenting were correct. I had collected them from raw data, personally, and I stood behind every figure. But the new leader was not engaging with the data. He was engaging with his own perspective on what the data should show. Every approval took longer than it should have.

Because of this, we lost so much time and effort and the project was delayed massively. Not because of a technical failure. Because of a leadership failure. A leader who filters information through ego does not improve their organization. They protect their self-image at the organization's expense.

MANAV CHADHA

Before any report drives a significant decision, ask: can we trace this number back to the raw data right now, in this room? If the answer is no, the number is not ready to drive a decision.

The AI Governance Maturity Model

Level	Description
Level 1: Unaware	AI tools deployed. No policy. No owner. No audit trail. Most organizations are here.
Level 2: Reactive	Policy exists on paper. Controls triggered only by incidents. Governance theater.
Level 3: Managed	Risk Register active. Vendors audited. Human review enforced. Data classification published.
Level 4: Proactive	Governance is a competitive advantage. Every principle operational, tested, and measurable.

WHAT TO DO NEXT

Before your next board presentation, spend one hour with the team that manages your most critical systems. Ask them: what is the biggest risk leadership does not know about?

Every governance dashboard metric should have a source field: where the number came from, when last verified, and who is responsible. Metrics without sources are decorations.

Build a thirty-day habit: ask your CISO or CTO each week for one piece of bad news. Respond with a question, not a judgment.

CHAPTER 5

Clean Fuel: Data Integrity

Garbage in, governance out. There is no shortcut past documentation.

Documentation is not bureaucracy. Documentation is survival.

The single most common root cause I find in organizations that have experienced a significant technology failure is not inadequate technology. It is the absence of documentation. The absence of a starting point. The absence of a record of what the organization had, how it worked, who was responsible for it.

When I started looking for evidence to support the CMMI assessment, I found something I was not fully prepared for. The absence of documentation was not a gap in one area. It was organizational. Systemic. Nobody had managed documentation consistently. Nobody had looked at the environment from the perspective of an auditor who had never been inside it before.

The answer, in most of the teams I assessed, was no. Nobody could tell me: if a new person joined this team tomorrow, could they understand what we have and what to do if something breaks?

That is frustrating. I will not pretend it was not. When you spend hours looking for a document that should have been created years ago, that would have taken one afternoon to write and would have saved days of reconstruction, there is a specific kind of frustration that comes from that. We were so frustrated going through those systems. Nobody collaborated well. Everyone was doing things the way they individually thought was right.

Organizations run on shared understanding. And shared understanding requires documentation.

FROM EXPERIENCE

The best organizations treat documentation as a first-class deliverable. Not something you do after the work is finished. Something you do as part of the work, because the work is not finished until anyone who needs to understand it can understand it from what is written down.

The Documentation Standard That Matters

Can a qualified person who has never worked in this organization, given nothing but the documentation, understand what the critical systems are, how they connect, what the correct procedure is for common failure scenarios, and who to call if the procedure does not resolve the issue? Most organizations cannot meet that standard. Setting it changes what documentation looks like. It is no longer a log of what was done. It is a guide

for what to do.

Store your most critical documentation in at least two locations: one online with version control, one offline accessible without network access. A crisis that takes your systems offline also takes your online documentation offline. That is exactly when you need it most.

WHAT TO DO NEXT

Identify your three most critical systems. Ask the person responsible for each: could your replacement understand this system from the documentation that exists? If the answer is no, schedule the documentation work before anything else.

Every project should have documentation as a required deliverable. If a project ends without documentation, it is not complete.

Store your most critical documentation both online with version control and offline. Test the offline version quarterly.

CHAPTER 6

The Third-Party Trap

Your vendor's breach is your breach. There is no legal distinction that matters to your clients.

We were in the middle of a global migration project. The vendor told us the software would support the migration completely. Full compatibility. Everything we needed them to do, they confirmed they could do.

The 32-Bit Architecture Red Flag

When I started to dig into the technical details, something did not add up. What eventually came out was that the software only supported 32-bit architecture.

That was a problem. Not a minor configuration issue. A fundamental architectural limitation that meant the software could not support the migration the way we needed it to. When I raised this, their response was essentially: yes, that is how it is. As if this was normal. As if it was reasonable to sell a global migration solution that could not support a modern architecture.

I never expected this. It was a big red flag. And according to them, it was all normal. That vendor was not upgrading their systems. The migration could not proceed without major changes on our side. This is the Halo Effect in practice: the vendor had a professional sales team and good marketing materials, and all of that created an assumption of technical currency that the product simply did not have.

Brand is not a security control. A professional sales presentation is not evidence of technical quality.

The Contract Clauses That Matter Most

Clause	Why It Cannot Be Negotiated Away
24-hour breach notification	Under most regulations you have 72 hours to report a breach. Your vendor must notify you within 24 hours to give you adequate response time.
Right to Audit	The right to conduct 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.

WHAT TO DO NEXT

Create a vendor security questionnaire and send it to your three most critical vendors this quarter. Require written responses from their security lead, not their account manager.

Add a 24-hour breach notification clause to every vendor contract that handles sensitive data. Track which vendors push back.

When evaluating a new vendor, add one hour of specific technical due diligence. The answers to specific technical questions matter more than the demo.

CHAPTER 7

The Human-in-the-Loop

The smartest algorithm in the world cannot feel that something is wrong.

I believe in AI. I use it. I advocate for it. I have built technical solutions on top of it. But one hundred percent dependency on any single system, including AI, is a risk without a mitigation.

Keep the humans sharp.

The Team Member Who Caught What the System Missed

I was working with a team member who was responsible for managing technical systems according to documented standard operating procedures. This team member had decided, without telling anyone, that they had a better approach. Instead of following the SOPs, they were using open-source tools and their own judgment. The systems appeared to be running normally. There was no alert, no flag, no monitoring tool that indicated anything was wrong.

What I noticed was behavioral. When I started asking questions about specific technical decisions, the answers were not quite right. The nervousness of someone giving you an explanation they made up rather than one they remember. So I pushed. What emerged was that the SOPs had not been followed for some time, and the open-source approach had introduced changes that could have caused significant damage to the production environment.

I did not stay quiet. I coached the team member. I explained why the SOPs existed and why following established procedures protects both the organization and the individual. I reported the situation to leadership. Not to punish. To ensure the risk was assessed and the environment was verified as stable.

FROM EXPERIENCE

Speaking up is not the same as reporting someone. Speaking up, when done with care and accountability, is what protects teams from preventable disasters. The culture that makes it easy to speak up early is the culture that avoids the hard conversations later.

The Intellectual Autopilot Risk

The greatest risk of AI assistance is not that the AI gets something wrong. It is that humans stop checking whether it did. Every AI-generated output must have a human reviewer who is capable of challenging it. Not a formality. A qualified person who can say: this does not match what I know, this needs to be verified before we act on it.

WHAT TO DO NEXT

Build a culture where questions are welcome. The team member who comes to you with a concern is doing their job correctly. Respond with curiosity, not frustration.

Review your SOPs annually. SOPs written for a previous state of the environment are not protection.

When someone on your team catches something the systems missed, recognize it publicly. That recognition invites more of the same behavior.

CHAPTER 8

The Ethics Moat

The cover-up always costs more than the truth. Without exception.

I could talk about my mother every day.

Not because the grief is still acute, though the loss is real. But because the lessons she embedded in me are ones I keep finding in my professional life, in situations she never could have imagined but would have recognized instantly.

What She Taught Without Teaching

My mother ran her own businesses. She built them from very little, in conditions that were not always favorable, against challenges that would have caused many people to take shortcuts or to compromise.

She did not take shortcuts. When things went wrong, and things always go wrong, her first question was not who is to blame. It was what happened and what can we learn. She did a version of root cause analysis before I knew what root cause analysis was. She understood intuitively that if you do not understand why something went wrong, it will go wrong again.

She trained the people around her with the lessons she learned. Not by lecture. By example, and by creating a culture where learning from mistakes was expected and repeating them was not acceptable. Fail forward, learn the lesson, train others, do not repeat it.

She was the best teacher I ever had. She was the best governance practitioner I ever knew. And she would have laughed at both of those descriptions.

> MANAV CHADHA
>
> *The most powerful governance framework I have ever encountered was not a compliance standard or a maturity model. It was my mother's approach to mistakes: find the root cause, learn the lesson so deeply that it changes your behavior, and never make the same mistake twice. Everything else is commentary.*

Building Psychological Safety Into Governance

A leader can have the highest personal standards of transparency and still preside over an organization where people hide problems every day. Because the culture is set not by the leader's personal standards but by how the leader responds to bad news.

If you respond to bad news with anger, the next person who has bad news will wait longer before bringing it to you. If you respond with blame, the person after them will fix the problem before you ever know it existed.

Build it by responding to bad news with two questions: what happened, and what do we need to do about it? Not who is responsible and what are the consequences. The first pair builds an Ethics Moat. The second pair destroys it, one conversation at a time.

WHAT TO DO NEXT

Do a root cause analysis on your last three significant incidents. Not to assign blame. To find where honest early disclosure would have reduced the impact.

When someone brings you bad news early, thank them before you ask any other question. That one response, consistently applied, changes what information reaches you.

Share the lessons from past incidents across the organization. Not the blame. The lesson. Organizations that share lessons learn faster than organizations that keep them in reports nobody reads.

THE GOVERNANCE CASCADE

From the CEO's Desk to the Front Line

How a mandate travels through an enterprise and where it breaks down.

In a small organization, governance is a conversation. The CEO decides. The team follows.

In an enterprise, governance is a system. And like every system, it has failure points. Between the CEO who owns the accountability and the frontline employee who approves an AI tool request, there may be five or six levels of leadership. Each level is a handoff. Each handoff is a place where the mandate can be diluted, deprioritized, or quietly ignored.

Level	Role	What They Own	Common Failure Mode
Level 1	CEO	Sets the non-negotiable standard. Signs the mandate. Chairs the quarterly review.	Sets mandate but treats it as a one-time act. Does not follow up.
Level 2	CIO / CTO	Translates CEO mandate into architecture and approved tool inventory.	Translates into technology without making it operational for business units.
Level 3	CISO / CRO	Validates controls match the risk appetite. Owns vendor audit process.	Softens bad news before it reaches Level 1.
Level 4	Operations / Divisions	Enforce the mandate within their teams. Named AI tool owners.	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. Every level of the cascade must translate the mandate into something the next level can actually act on. When a level cannot translate it, the mandate stops there.

The Test That Reveals the Gap

Here is the question I ask every enterprise CEO.

"Without calling your IT team, right now, 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. That gap is fixable in ninety days.

CHAPTER 9

The Kill Switch

There is no kill switch. There is only preparation.

Every leader I have ever worked with has asked the same question in some form: if something goes catastrophically wrong with our AI or our technology systems, can we just shut it down?

The answer is more complicated than yes.

The First Sixty Minutes: What Nobody Tells You

I have been inside real technology crises. Not the theoretical versions described in tabletop exercises. The real ones, with incomplete information and no clean options.

I was in the same boat once. Scrambling here and there rather than approaching the right person to ask for the information needed to fix the crisis. What I wish someone had told me before it started: if there is documentation, if there is a place where I can find the information, that would have made life so much easier. Most of the time it is about finding things that are in the right place but you are looking in the wrong place under stress and fear.

> FROM THE CRISIS ROOM
>
> *The person who calms down first in a technology crisis is the most valuable person in the room. Not the most technically skilled. Not the most senior. The calmest. Calm allows thinking. Thinking finds the answer.*

Attitude Over Aptitude

In a technology crisis, attitude matters more than aptitude. No matter how smart you are: if you panic and shake every time a technology crisis happens and cannot speak up, cannot ask, cannot use your brain, then the aptitude is of no value.

There should not be any blame game. There should not be any throwing things into other people's courts. It is a team effort. Attitude matters a lot. That is not a soft concept. It is a governance infrastructure requirement.

What Preparation Actually Looks Like

After one significant technology crisis, the leader I was working with changed their entire approach to technology documentation. They asked the technology team to build documentation for everything that talked to everything else. Network diagrams. System interdependencies. Contact lists for every critical vendor.

Recovery procedures for every critical scenario. Version controlled. Updated continuously. Stored both online and offline.

No matter how big the disaster is, we should be able to use risk mitigation efforts immediately. If we cannot mitigate it, we should be able to avoid or transfer it. There is always a residual risk we have to accept. That is continuous risk management. Not a one-off exercise.

WHAT TO DO NEXT

Schedule a tabletop exercise this quarter that simulates a complete system outage. The gaps it reveals are your preparation roadmap.

Build and maintain an offline incident response package: network diagrams, key contact numbers, vendor support lines, step-by-step recovery procedures for your most critical systems.

Practice asking for help. In high-pressure situations the instinct is to keep trying alone. Build a culture where asking for help is a sign of intelligence, not inadequacy.

CHAPTER 10

The Future-Proof CEO

Most executives think AI governance is about technology. It is about behavior, culture, and the decision to lead.

I want to tell you what I actually hope for when someone finishes this book.

Not that they implement every framework in the appendix. I hope they become the leader who asks before the crisis. The executive who looks at a new AI tool proposal and says: before we deploy this, I want to understand how it handles our data and who is accountable if something goes wrong.

That question, asked consistently, changes everything downstream.

The Leader Who Changed

A leader I worked with went through a significant technology crisis with our team. Before the crisis, their view of technology risk was essentially the standard one: it is being managed by the people I pay to manage it.

During the crisis, they watched what happens when the people managing the technology are scrambling for information that was never captured. They understood something that no briefing had communicated. They understood that the gap between the technology team and the business leadership was a risk. A real one, with real consequences.

In the weeks after, they drove the process themselves. They asked the technology team to build documentation for everything. Version controlled, stored in multiple locations, tested regularly. They established a governance calendar. And they told me something I think about often: I used to think my job was to set the direction and let the team manage the details. Now I understand that some details are not details. They are load-bearing. And I cannot know which ones until I understand enough to ask.

MANAV CHADHA

You do not need to know how every system in your organization works. You need to know which ones would stop the business if they failed, who is accountable for each one, and what the recovery plan looks like. That is not a technical requirement. It is a leadership requirement.

The Non-Negotiables: For the CEO Who Is Not Technical

A fair question deserves a direct answer: how is a leader without a technology background supposed to know which technical issues are non-negotiable and which are ordinary operational noise? You cannot evaluate the technology. You do not need to. You evaluate the answers. The following five conditions are non-negotiable

in any organization, and recognizing their absence requires no technical knowledge at all.

Non-Negotiable	How You Recognize a Violation Without Technical Knowledge
Every critical system has a named owner.	You ask who owns it and receive a department name, a vendor name, or a pause instead of a person's name.
The Risk Register can be produced on demand.	Someone needs time to compile it. A register that requires assembly is a filing project, not a control.
Recovery has been tested, not just planned.	Nobody can tell you the date of the last test or what changed because of it.
No AI tool touches your data without written zero-retention proof.	The answer references a privacy policy or a verbal assurance instead of a contract clause.
Bad news travels up faster than good news.	You learn about problems from incidents instead of from reports. Every crisis that surprises you is a violation of this one.

When any of these five is violated, you do not need to understand the underlying technology to act. You need to ask the question again, in writing, with a deadline, and keep asking until the answer is a name, a date, or a document. That persistence is the entire technical skill the job requires of you.

The Monday Morning Diagnostic

Question	What a No Tells You
Can someone produce the Risk Register in five minutes?	Governance is theoretical. The Register exists as a project, not a tool.
When was the last tabletop exercise on the BCP?	The BCP is a document, not a plan. It has not been tested against reality.
What percentage of employees could pass an AI data handling test today?	The policy exists. The behavior it requires does not.
When did you last exercise the Right to Audit with your most critical vendor?	Vendor governance is contractual, not operational.
If your primary AI tool went offline right now, how long until operations are impaired?	You have built AI dependency. That is a different and more dangerous thing than AI assistance.

The Last Word

I want to tell you something I have not said anywhere else in this book.

I do not know how many people will read this. I do not know whether it will reach a hundred people or a hundred thousand. I have spent 25 years in IT and over a decade in cybersecurity, and the honest truth is that I did not write this book because I had a marketing strategy. I wrote it because I kept walking into organizations

where the gap between what leadership believed was happening and what was actually happening was costing real people real consequences, and nobody had written the thing I wanted to hand them.

So I wrote it.

My mother passed away on June 12th, 2025. I am publishing this book on June 12th, 2026. That date is not a coincidence. It is an intention.

She never read a governance framework in her life. She never sat in a boardroom. She never attended a security conference or signed an audit report or reviewed a Risk Register. But she understood something that I have spent 25 years trying to put into professional language. She understood that when you are in a position of responsibility, you do not get to decide which consequences belong to you and which ones belong to someone else. You own them. All of them. Before the crisis reveals them and after. That is not a burden. That is the job.

Every principle in this book, I learned it first from watching her. The root cause analysis before I knew the term. The fail forward approach before I had the words. The insistence on doing the right thing quietly, without an audience, when the easier option was right there in front of her. She built businesses the same way she raised a family. With accountability that did not need an external audit to function. With transparency that did not need a policy to enforce it. With integrity that did not need a consequence to motivate it.

That is the Ethics Moat. That is the Governance Cascade working at the most personal level. That is the mandate that lives not on a signed declaration page but in the character of the person who leads.

To the CEO who has read this far: you already know what you need to do. You knew before you opened this book. What you needed was someone to tell you that it is your job, not the IT department's job, and that the cost of not doing it is, in nearly every case I have seen, higher than the cost of doing it now.

That is what I needed someone to tell me. That is what my mother showed me.

Go build the governance your organization needs. Not because a regulator is watching. Not because a consultant told you to. Because the people downstream of your decisions deserve a leader who owns the accountability completely.

That is the mandate.

That is all it has ever been.

THE CORE TRUTH

> *Most executives think AI governance is about technology. The truth is, it is about behavior. How responsibly we use the tools available to us. Whether we let them become dependencies we cannot explain. Whether we take shared accountability for what they produce. That is the choice this book is asking you to make. Manav Chadha | The AI Mandate | June 12, 2026*

WHAT TO DO NEXT

Print the Monday Morning Diagnostic. Answer it honestly, without preparation, right now. Circle the questions you could not answer confidently. Those are your governance gaps.

Identify one person in your organization who will own the governance function going forward. Give them the title, the budget, and direct access to you.

Put the 90-Day Sprint start date in your calendar today. The sprint that starts in three months does not start in three months. It starts never.

Book thirty minutes with your technology leadership team this month. Not to review a dashboard. To ask: what is the governance gap you have been most concerned about that you have not yet found the right way to bring to me?

THE 90-DAY GOVERNANCE SPRINT

From Reading to Running

The sprint below is a compressed, prioritized sequence of the most important governance actions in this book. It is designed for the leader who wants to go from zero formal AI governance to a defensible, operational governance posture in ninety days.

Phase	Days	Priority Actions
Phase 1 Foundation	1 to 30	Produce the Risk Register. Audit every AI tool in use. Assign named owners. Run BCP tabletop. Issue CEO mandate to all staff.
Phase 2 Vendor and Data	31 to 60	Send Vendor AI Audit Questionnaire. Read SOC 2 Type 2 reports. Verify zero retention clauses. Publish Data Classification Policy. Run staff training.
Phase 3 Policy and Monitor	61 to 90	Publish AI Acceptable Use Policy. Name Governance Monitor. Run AI tabletop. Hold first formal Governance Review. Lock quarterly calendar.

MANAV CHADHA

The sprint that starts in three months does not start in three months. It starts never. Pick a start date. Put it in the calendar. The best governance program in the world is the one that actually begins.

APPENDIX A

IT Risk Assessment Template

Know what you have. Know who owns it. Know what happens if it fails.

Run this exercise in Week 1. Survey every department head. Assign a named Business Owner and Technical Owner to every entry. An AI tool with no named owner has no governance.

AI Tool Inventory

Tool Name	Dept	Approved	Business Owner	Tech Owner	Risk	Last Reviewed
		Y/N			H/M/L	
		Y/N			H/M/L	
		Y/N			H/M/L	
		Y/N			H/M/L	

Critical System Register

System	Function	Named Owner	RPO (hrs)	RTO (hrs)	Risk Notes

Risk Rating Guide

Rating	Definition	Required Action
HIGH	Loss would cause immediate operational, financial, or regulatory harm.	CEO notified within 24 hours. Quarterly audit required.
MEDIUM	Loss would cause significant disruption but organization can operate in degraded mode.	CISO/CTO notified within 48 hours. Semi-annual review.
LOW	Loss would cause inconvenience but minimal business impact.	Standard incident process. Annual review sufficient.

APPENDIX B

Vendor AI Audit Questionnaire

A vendor who hesitates to answer how they protect your data, that hesitation is their answer.

Send this to every vendor that processes your organizational data. Require written responses from their security lead. Not their account manager. Give them ten business days.

Question	Vendor Response
Do you retain customer prompts or data submitted to your AI tools?	
Can you provide your current SOC 2 Type 2 report?	
Have you experienced any data breach in the last 24 months?	
What is your breach notification timeline?	
Where is our data stored geographically?	
Do you use our data to train or improve your AI models?	
What encryption standard do you apply to data in transit and at rest?	
Do you conduct annual penetration testing?	
Who has access to our data within your organization?	
Do you share our data with any third parties or sub-processors?	
How do you handle deletion of all our data upon request?	
What certifications does your organization hold?	
Are you willing to include a 24-hour breach notification clause in our contract?	

Scoring Guide

Finding	Risk	Action
Cannot produce SOC 2 Type 2	HIGH	Do not approve until report is provided.
Retains prompts for model training with no opt-out	HIGH	Require contractual opt-out or reject.
Breach notification exceeds 72 hours	HIGH	Negotiate 24-hour clause or reject.
Data stored outside approved jurisdiction	MEDIUM	Require contractual data residency.

Finding	Risk	Action
No annual penetration testing	MEDIUM	Require testing schedule as condition.
All questions answered clearly	LOW RISK	Proceed to contract review.

APPENDIX C

AI Acceptable Use Policy

A rule that comes from IT is optional. A rule that comes from the CEO is policy.

This policy must be signed and issued by the CEO. Policies that come from IT are treated as IT's problem. Policies that come from the CEO are treated as everyone's responsibility.

1. Purpose

This policy establishes the rules governing how employees of [Organization Name] may use artificial intelligence tools in the course of their work. It applies to every employee, contractor, and consultant, regardless of role, location, or seniority.

2. The Three Rules

Rule	Requirement
Rule 1: Public Domain Standard	If you would not post this information publicly, you do not feed it to a public AI tool. Includes client names, financial data, personnel information, internal project details, and any proprietary business information.
Rule 2: The Attachment Ban	No internal documents, financial reports, client files, personnel records, or proprietary content may be uploaded to any public AI tool under any circumstances.
Rule 3: Zero Retention Rule	Before using any AI tool for organizational purposes, verify that the vendor does not retain your prompts or data. If unsure, do not use the tool. Refer to the Approved AI Tools list.

3. Consequences

Breach of this policy may result in formal disciplinary action up to and including termination. In cases involving regulatory breach, personal liability may also apply.

CEO Signature: _____ Date: _____

APPENDIX D

Data Classification Policy

You cannot protect what you have not classified.

Four tiers. Plain language. Any employee should be able to classify their own data without calling IT. When in doubt, classify higher.

Tier	Name	Definition	Examples	AI Tool Rule
1	Restricted	Highest sensitivity. Loss would cause severe regulatory, financial, or reputational harm.	Health information, payment card data, legal privilege, M&A; details	Never enter into any AI tool.
2	Confidential	Sensitive business information. Loss would cause significant competitive harm.	Client contracts, financial projections, personnel files, internal pricing	Approved internal tools only, with explicit permission.
3	Internal	General business information not intended for public view.	Internal policies, meeting notes, project plans, org charts	May be used with approved AI tools.
4	Public	Information approved for public release.	Press releases, published reports, public job postings	May be used with any approved AI tool.

APPENDIX E

Quarterly AI Governance Review Agenda

A governance meeting that does not produce decisions is a status update. This is not that.

This meeting happens once per quarter. The CEO chairs it. Every member of the senior leadership team attends. Minutes are formally recorded. Every action item has a named owner and a deadline. Total time: 90 minutes.

Time	Agenda Item	Owner	Expected Output
0:00 to 0:10	Risk Register Review: what changed since last quarter, any new HIGH items.	CISO / CRO	Updated register confirmed or escalated.
0:10 to 0:25	AI Tool Inventory Update: new tools added, unapproved tools discovered, owners confirmed.	CIO / CTO	Clean approved tool list with named owners.
0:25 to 0:40	Vendor Status: SOC 2 reports reviewed, breach notifications received, contract renewals.	CISO	Vendor risk status confirmed or flagged.
0:40 to 0:55	Policy Compliance: AUP adherence, training completion rates, any policy breaches.	CIO / HR	Compliance rate reported, gaps actioned.
0:55 to 1:10	Incident Review: any incidents since last quarter, root cause, what changed.	CISO	Incidents closed or escalated with named fix owner.
1:10 to 1:25	Open Issues from Last Quarter: were all action items completed?	All owners	Outstanding items escalated to CEO if overdue.
1:25 to 1:30	New Action Items: decisions made today, named owners, deadlines.	CEO	All new items recorded in minutes before meeting closes.

CEO Signature confirming minutes are accurate: _____ Date: _____

APPENDIX F

Incident Response Checklist

Print this. Keep it somewhere accessible. The day you need it is not the day to look for it.

This checklist covers the first 72 hours of a technology incident. The sequence matters. Do not skip steps. Do not reorder them.

First 60 Minutes: Stop the Propagation

Step	Action	Owner	Done
1	Confirm the incident is real. Get one named technical person to confirm the nature and scope.	Technical Lead	[ ]
2	Activate your incident response team. Named individuals only.	CEO / CISO	[ ]
3	Isolate affected systems if possible. Goal is to stop spread, not to fix the problem yet.	Technical Lead	[ ]
4	Preserve evidence. Do not wipe or restore anything until forensic capture is confirmed.	Technical Lead	[ ]
5	Identify what data may be affected: client, employee, or financial data.	CISO	[ ]
6	Notify the CEO directly: what happened, what is affected, what we are doing right now.	CISO	[ ]

Hours 1 to 24: Assess and Contain

Step	Action	Owner	Done
7	Determine blast radius: what systems, what data, how many people affected?	CISO / CTO	[ ]
8	Check vendor notification obligations.	CISO	[ ]
9	Notify legal counsel before any external communication.	CEO	[ ]
10	Review your regulatory notification obligations and timeline.	Legal / CISO	[ ]
11	Draft initial internal communication to affected staff. Factual only. No speculation.	CEO / Legal	[ ]
12	Continue containment. Executives remove obstacles, do not direct the technical team.	Technical Lead	[ ]

Hours 24 to 72: Notify and Begin Recovery

Step	Action	Owner	Done
13	File regulatory notifications if required.	Legal / CISO	[ ]
14	Notify affected clients if their data was involved. Be honest. Be specific.	CEO	[ ]
15	Notify your board chair. One page: what happened, what data, what we are doing.	CEO	[ ]
16	Begin documented root cause analysis.	CISO / CTO	[ ]
17	Begin documented recovery. Every step recorded. Nothing done from memory.	Technical Lead	[ ]
18	Establish daily CEO briefing until incident is closed.	CISO	[ ]

After Recovery: Root Cause and Prevention

Step	Action	Owner	Done
19	Formal post-incident review within 14 days. Not a blame session. A learning session.	CEO chairs	[ ]
20	Update the Risk Register to reflect what this incident revealed.	CISO	[ ]
21	Update the AI Acceptable Use Policy if any AI tool was involved.	CIO / CISO	[ ]
22	Communicate lessons learned to all staff. Honest, brief, and useful.	CEO	[ ]
23	Confirm with Board that incident is closed and prevention measures are in place.	CEO	[ ]

Incident Commander: _____ Opened: _____ Closed: _____

About the Author

Manav Chadha is a technology leader, cybersecurity practitioner, and governance advisor with 25 years of experience in IT, information security, and cybersecurity.

He has worked across the energy sector, healthcare, municipal government, retail, and professional services. He has led enterprise governance programs, conducted high-stakes international audits, and managed real technology crises from inside the organizations experiencing them.

He developed a semantic leakage prevention prototype after observing how confidential organizational data was entering public AI tools without governance, oversight, or the knowledge of the leaders accountable for it.

He holds a Master of Business Administration (MBA), Certified in Cybersecurity (CC) from ISC2, Certified Information Security Manager (CISM) from ISACA, Associate C|CISO from EC-Council, EC-Council Certified Security Analyst (ECSA), Certified Cybersecurity Educator Professional from Red Team Leaders, ISO 27001 Lead Auditor and ISO 42001 Lead Auditor from Exemplar Global.

He has mentored more than two hundred professionals in IT, cybersecurity, and governance over the course of his career.

He runs his own advisory practice through Quality Training and Technology Services Inc., based in Edmonton, Alberta, Canada, helping organizations build practical, honest governance postures.

He lost his mother in 2025. This book is dedicated to her, to his wife and children, and to every leader willing to be accountable before the crisis demands it.