BOARD DEFENSIBLE

The CISO's 12-Step Program for Building an AI Security & Risk Program

Chapter 2

Not All AI Is Equal

Pillar 2 — AI Classification & Risk Tiering

PART ONE — FOR THE CISO

The Executive Briefing

Read this in the next ten minutes. It's what you need to walk into a board meeting, staff a workstream, and know whether your program is real or performed.

The Reality Check

Every AI security conversation starts the same way: someone asks "how many AI tools do we have?" and the honest answer is "more than we think, and fewer than we'll admit to the board." That's not a discovery problem — it's a categorization problem. Most organizations are trying to govern AI with a single policy, a single risk questionnaire, and a single approval workflow, as if the writing-assistant plugin in Legal's browser carries the same risk profile as an autonomous agent with write access to your CRM.

It doesn't. And treating it like it does is why AI governance programs stall. Some AI is invisible and low-risk, already embedded in tools you've had for years. Some is autonomous and high-risk, capable of taking actions in your environment without a human in the loop. Your governance program has to flex across that entire range, or it will either strangle the low-risk work or wave through the high-risk work. Both failures land on your desk.

Why This Matters at the Board Level

Boards don't want an inventory. They want to know whether the organization is exposed, where, and how much it would cost if something went wrong. A flat list of "47 AI tools in use" tells them nothing — it's noise dressed up as diligence. A tiered risk view tells them: 40 of those tools carry negligible incremental risk, 5 need contractual and monitoring controls, and 2 require board-level sign-off because they can act autonomously on sensitive systems.

That second version is a governance program. The first version is a spreadsheet. Audit committees are asking sharper questions about AI risk classification specifically because "we have an AI policy" stopped being a satisfying answer sometime in the last eighteen months.

The Board Conversation

Three things to remember every time you're in that room:

- **Lead with tiers, not tool counts.** "We've classified our AI footprint into four risk tiers, and here's where the real exposure concentrates" beats any number of tools discovered.
- **Name the autonomous systems specifically.** Boards understand that a chatbot and an agent with transaction authority are different categories of risk. Don't let them get flattened into "AI" as a single line item.
- **Show the trend line, not just the snapshot.** The tier distribution today matters less than whether Tier 4 (autonomous/agentic) is growing quarter over quarter — that's the number that should set the pace of your program.

Business Translation for the Board

> ***We can't put a dollar figure on 'AI risk' as a category, but we can tell you exactly what it costs to leave our two highest-risk AI systems ungoverned — and it's a number worth funding a program to avoid.***

Bringing In CTO, CIO & Business Teams

This pillar fails if it stays a security-only project. Use this to make the ask concrete when you're recruiting each function.

Stakeholder	Why this affects them	What we need from them
CTO / Engineering	Owns the technical architecture AI systems plug into, including everything built in-house.	A named engineering point of contact, and tiering built into architecture review before deployment, not after.
CIO / IT leadership	Owns the SaaS portfolio and the vendor contracts where embedded AI features quietly appear months after signing.	A 'has this vendor shipped new AI features?' question added to every renewal, looped into IT asset management.
Business unit leads	Their teams are the ones actually adopting AI tools to move faster — usually without asking anyone first.	Completion of the quarterly self-certification survey, and a named owner for every AI tool their team uses.
Legal / Privacy	Owns contractual and regulatory exposure across every tier.	Input on which tiers trigger DPA / vendor addendum requirements.

Stakeholder	Why this affects them	What we need from them
Data / AI / ML teams	Own the models the organization builds in-house — the systems most likely to be excluded from vendor-focused reviews.	Classification of internally built models alongside third-party ones, as a required SDLC field.

Program Objectives

Four things this pillar has to deliver before anything downstream — access, monitoring, incident response — can be considered real.



The Roadmap

This is the pace CISO Next Gen members have used to move from an ungoverned AI footprint to a board-reportable tiering program — without stalling the rest of the security program to do it.

BOARD DEFENSIBLE · PILLAR 2

Roadmap: Classify & Govern

From first tiering pass to continuous board reporting



This week's action: pick five AI tools already in use and score them against the four-tier framework.

Source: Board Defensible — Chapter 2 · CISO Next Gen

Metrics to Track

Five is the ceiling before a board stops actually tracking them. Report these on the same cadence every quarter — consistency matters more than volume.

Metric	What it measures	Target direction
Tier coverage rate	% of known AI systems classified into a tier, against the total identified across discovery.	Rising toward 100%
Tier 4 (see risk tiering template) count & trend	Raw count of autonomous/agentic systems, tracked quarter over quarter.	Watched closely, not necessarily falling
Time to classification	Median time between an AI tool's discovery and its tier assignment.	Falling toward days, not weeks
Intake vs. discovered	% of new AI adoption that went through intake vs. was found after the fact.	Rising toward 100% intake
Compensating control rate	% of Tier 3–4 systems with documented compensating controls.	Rising toward 100%

PART TWO — FOR YOUR TEAM

The Implementation Playbook

This is the working section — hand it to whoever owns execution. It defines the four categories of AI assets, what a complete tier entry looks like, the 30/90/365-day build sequence, and the failure modes that derail this work most often.

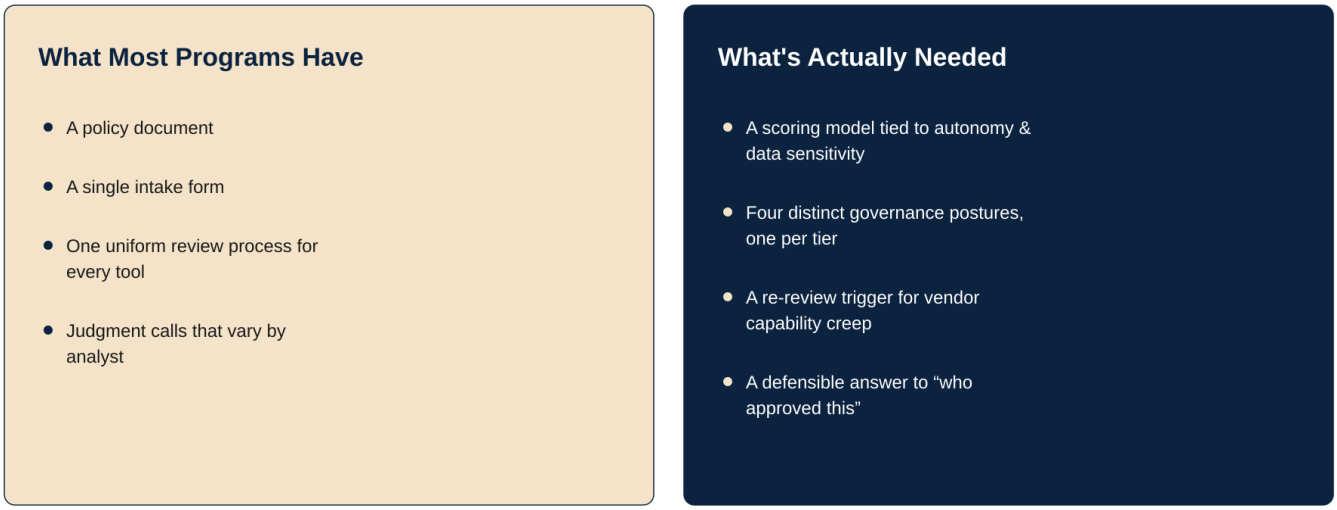
The Governance Gap

Before tiering anything, your team needs to internalize why most programs stall. It isn't a missing policy — it's a missing scoring layer between the policy and the decision.

BOARD DEFENSIBLE · PILLAR 2

The Governance Gap

What most AI programs have, versus what a defensible tiering program needs



Source: Board Defensible — Chapter 2 · CISO Next Gen

The Four Asset Categories

Every AI system your organization touches falls into one of four tiers, based on two variables: autonomy (does it act, or does it inform a human who acts?) and data sensitivity (what does it touch?).

1. Embedded & Invisible

AI features baked into tools you already use and have likely never separately reviewed: writing assistants, meeting transcription, smart-compose features, recommendation engines inside SaaS platforms. A meeting transcription feature that summarizes an internal call is Tier 1; the same feature auto-emailing that summary to an external distribution list has quietly become something closer to Tier 3.

Governance posture: *Light-touch. Register it, confirm it inherits the underlying vendor's existing contract terms, move on.*

2. Predictive & Analytical

Traditional ML and statistical models used for forecasting, scoring, or classification — fraud detection, credit risk scoring, demand forecasting, churn prediction. Informs human decisions; doesn't act on its own. Often touches sensitive or regulated data, which is why this tier existed under the name "model risk management" long before "AI governance" became a category.

Governance posture: *Standard model risk management — validation, bias testing, documented data lineage, periodic re-validation.*

3. Generative & Conversational

LLM-based chatbots, copilots, and content-generation tools, whether built in-house or purchased. No autonomous multi-step action, but capable of processing and potentially exposing sensitive inputs, and capable of producing outputs that get acted on directly by end users. This is the tier most existing policies were actually written for — which is part of the problem, since it's also the tier most organizations have already moved past.

Governance posture: *Data handling review, output monitoring, clear usage policy, DLP controls on inputs.*

4. Agentic & Autonomous

Systems that take multi-step actions with minimal human review — agents with tool access, systems connected via MCP or similar protocols, anything with write access to production systems, financial systems, or customer data. This tier didn't meaningfully exist in most enterprises three years ago and is now the fastest-growing category by adoption rate.

Governance posture: *Full pre-production review, least-privilege access design, human-in-the-loop checkpoints for consequential actions, continuous monitoring, and — increasingly — a named executive owner per agent, not per department.*

The tier boundary that matters most right now is the line between Tier 3 and Tier 4. A chatbot that drafts an email is fundamentally different from an agent that sends it, and most existing AI policies were written before that distinction needed to exist.

From the Field: Real CISO Stories

These are composite scenarios drawn from patterns CISO Next Gen members report repeatedly — not any single named organization, but the kind of story every practitioner in the room recognizes immediately.

The forecasting tool that upgraded itself

A CISO at a global pharmaceutical distributor discovered, six months into a formal AI inventory effort, that a demand-forecasting tool procured by the supply chain team two years earlier had quietly been upgraded by its vendor to include an 'autonomous reorder' feature — turning a Tier 2 tool into a Tier 4 tool without a single internal conversation. The contract hadn't changed. The risk profile had, completely. The team only caught it because they'd started asking 'has this tool's capability changed' as a standing question, not just 'do we still use this tool.'

The self-certification bet that paid off

A CISO at an asset management firm built their tiering model around a simple rule that made adoption easy: any team could self-certify a tool as Tier 1 using a five-question form, no security review required, as long as they were willing to have it audited quarterly. They got radical honesty because they made the cheap path the honest path — the moment self-declaring became easier than hiding it, teams stopped playing whack-a-mole with their own security function.

The agent nobody in engineering wanted tiered

A CISO at an HR technology platform found the hardest part wasn't building the framework — it was getting engineering to apply it to systems they'd built internally, not just vendor tools they'd bought. Everyone wanted to point the risk framework at procurement; nobody wanted to point it at their own roadmap. Tiering only stuck once it became a required field in the team's own SDLC ticket templates.

The agent hiding inside a statement of work

A CISO at a regional financial services firm had an internal audit surface a Tier 4 candidate nobody in security had reviewed: a customer service agent, built by a vendor's professional services team during an implementation project, quietly given write access to account notes and case status. It had never appeared in any AI inventory because it wasn't purchased as an 'AI feature' — it was purchased as 'implementation services.' Every statement of work involving a vendor's professional services team now requires an explicit AI-capability disclosure clause, reviewed before the SOW is signed.

> ***The pattern across all four: the gap wasn't a missing tool. It was a missing re-review trigger, a missing internal owner, or a missing line item nobody's inventory process was built to catch.***

How to Build It: Manual and Automated Methods

Most teams default to one approach and miss what the other catches. Manual methods surface intent and ownership no tool can infer. Automated methods surface scale and blind spots no interview will uncover. A real tiering program uses both, run in parallel, reconciled against each other on a set cadence.

Manual methods — run these first

- Five-question intake form: scores autonomy and data sensitivity against the tier rubric in under five minutes per tool.
- Business unit self-attestation survey: a short, recurring form asking each team to list AI tools in active use — low completion rates are themselves a signal.
- Contract and renewal audit: pull every active SaaS contract and renewal due in the next 12 months; ask each vendor directly whether AI features have been added since signing.
- Structured changelog review: assign someone to read release notes for your top 20 SaaS vendors monthly — this is often the only place capability changes are disclosed before a user tells you.

Automated methods — run these continuously

- Browser telemetry / DLP: tune existing DLP or endpoint tooling to flag traffic to known AI domains — the fastest way to surface unsanctioned personal-account use.
- CASB / SSPM platforms: discover connected and OAuth-authorized apps automatically, including AI features employees have granted access to company data.
- API and network egress monitoring: flag outbound calls to LLM provider APIs directly from your network — catches engineering-built integrations that never went through procurement.
- Identity and access log review: query your IdP for OAuth grants and service accounts created in the last 90 days — new agent and integration accounts often appear here before anyone reports them.

> ***Reconciliation is the step teams skip. Once a month, compare what your automated tools found against what your manual tiering has logged. Every unmatched entry on either side is a gap — either a tool no one interviewed about, or a tool your scanning missed.***

Implementation Guide — 30 / 90 / 365

The 30-day foundation

Week 1 is building and pressure-testing the tiering rubric with 10–15 known AI systems you can already name. Weeks 2–4 focus on securing business unit and engineering sign-off on the tier definitions — arguing over definitions later, one tool at a time, is a much slower fight.

The 90-day build-out

Run full discovery across procurement records, SSO logs, and business unit self-reporting. Classify everything found. Stand up the standing review board for Tier 3–4 systems. Report the first tier distribution to the board.

The 12-month target

Intake becomes the default path for all new AI adoption, not an afterthought. The Tier 1–2 audit program runs on autopilot. Every Tier 4 system has a named executive owner and documented compensating controls. Quarterly board reporting on tier distribution and the Tier 4 trend line is routine.

Common Failure Modes

- Building the framework before securing the mandate to enforce it — a tiering model without teeth becomes a taxonomy exercise nobody uses.
- Making every tier feel like Tier 4 — if self-certifying a writing assistant takes the same effort as reviewing an autonomous agent, teams stop disclosing anything.
- Treating classification as a one-time event — vendor capability creep is one of the most common ways a tool silently moves up a tier without anyone noticing.
- Excluding internally built systems — the framework has to apply to what your own engineering team ships, not just what procurement buys.
- Excluding AI bundled into professional services — vendor-delivered AI capability doesn't always show up as a line item; sometimes it shows up inside a statement of work.
- No named owner for Tier 4 systems — autonomous systems without a specific accountable executive are the ones that cause board-level incidents.

This Week's Action

> ***Pick five AI tools your organization already uses — pull them from procurement records, not memory — and manually score them against the four-tier framework above. You will almost certainly find at least one that's a tier higher than whoever approved it assumed.***

Chapter Toolkit

This chapter includes the AI Risk Tiering Scorecard (five-question intake form), a Vendor Capability-Change Questionnaire for contract renewals, a Tier 3–4 Review Board Charter template,

a Board Reporting one-pager, and sample SDLC ticket-field language for internally built systems. Provided alongside this chapter.

Exclusive Invitation for CISO NextGen Members:

If you are a CISO or Security leader within an enterprise, please reach out for an invite to our first CISO NextGen NYC Dinner Summit: Inside Mythos: GlassWing Lessons for the AI Era

📍 New York, NY · September 29 5:30-8:30pm

A closed-door mini-summit and dinner for security leaders navigating the next wave of AI risk. Security Executive from a leading bank who's leading GlassWing remediation efforts, opens with hard-won lessons from the front lines. Use these lessons to defend your own organization.