BOARD DEFENSIBLE

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

Chapter 1

# You Can't Govern What You Can't See

*Pillar 1 — AI Inventory & Visibility*

---

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

Ask any CISO a simple question — "do you have a complete list of every AI system running in your environment right now?" — and watch the pause before the answer. The honest response is almost always some version of no.

That's not a failure of leadership. AI entered your enterprise through a hundred doors at once: procurement, engineering sprints, vendor product updates nobody flagged, and thousands of individual employee decisions made without asking anyone. No single team owns all of those doors. That's the problem this chapter solves.

### Why This Matters at the Board Level

Every AI security program — inventory, access, monitoring, incident response — depends on one prerequisite: knowing what exists. A board that hears "we've deployed AI governance controls" without first hearing "and here is our complete inventory" is hearing a program built on sand. Regulators are converging on the same expectation. You cannot attest to what you cannot enumerate.

### The Board Conversation

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

- They're focused on business risk, not technical risk. Translate every finding into money, operations, or reputation — and stop there. “We found 40 shadow AI tools” means nothing; “we closed exposure that could have cost us a regulatory fine” does.
- They have minutes, not hours. Lead with the one sentence that matters — the roadmap stage, the metric trend, the ask. If they want the detail behind it, they'll ask; don't front-load it and bury your point.
- They're judging the trend, not the moment. A single number means nothing without direction. Never show a metric without also showing which way it's moving and what you're doing about it.

Frame this pillar in outcomes, not activity. Don't say “we're building an AI inventory tool.” Say: “We now know what AI is running across the organization, who owns each system, and what data it touches — and we review that list on a defined cadence.” Show the roadmap and the metrics later in this section, not the raw asset count — a big number without structure reads as noise, not maturity. Never present the inventory as finished; present it as governed, current, and measured.

## Business Translation for the Board

> ***This inventory doesn't just reduce risk — it pays for itself by surfacing the unapproved, duplicated, and forgotten AI subscriptions quietly draining budget across every team, often covering the cost of the program before you count a single averted breach.***

## 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** | Most shadow AI and every custom-built agent originates in engineering — often without a formal security review before it ships. | A named engineering point of contact for the inventory, and an AI/agent flag 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. | An ‘has this vendor shipped new AI features?’ question added to every renewal, and IT asset management looped into discovery. |
| **Business unit leaders** | Their teams are the ones actually adopting shadow AI tools to move | Completion of the quarterly self-attestation survey, and a named |

| Stakeholder | Why this affects them | What we need from them |
|---|---|---|
| | faster — usually without asking anyone first. | business owner for every AI tool their team uses. |
| **Procurement** | Sees every new tool request before security does, and sees every contract renewal first. | An AI-disclosure question added to intake and renewal workflows, flagged automatically to security. |
| **Finance** | Sees shadow AI spend before security does — personal-card subscriptions show up in expense reports first. | A standing flag on recurring charges to known AI vendors during routine expense review. |

## Program Objectives

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

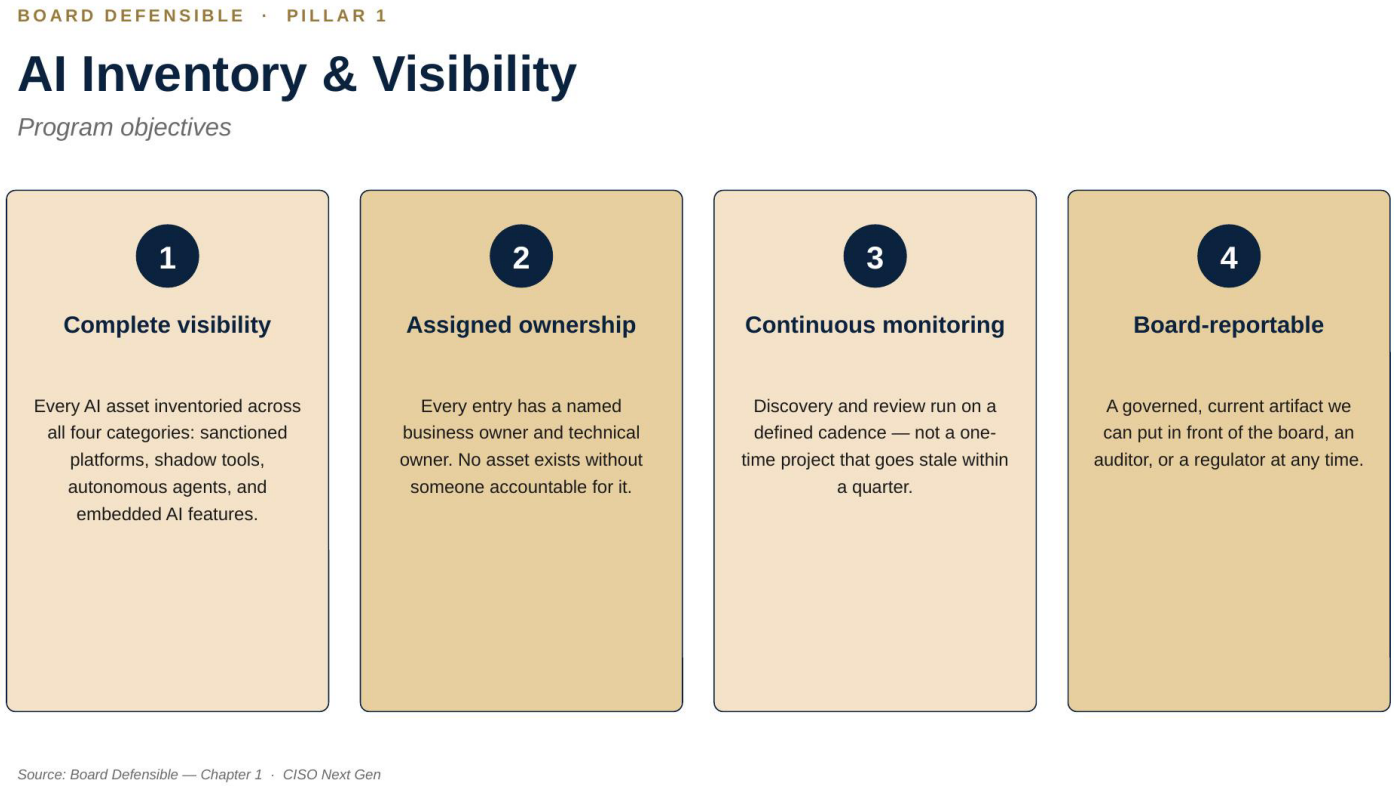

## The Roadmap

This is the pace CISO Next Gen members have used to move from partial visibility to a governed, board-reportable inventory — without stalling the rest of the security program to do it.

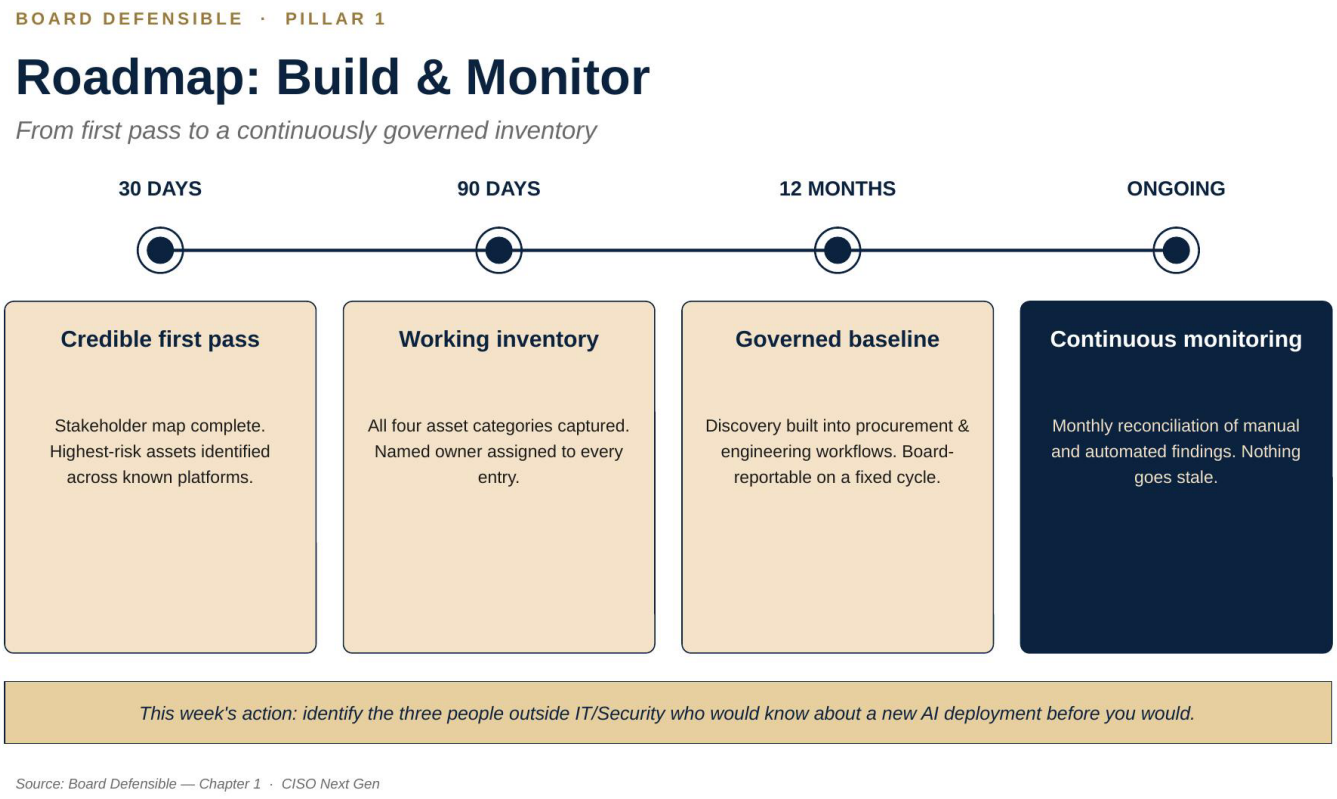


## Metrics to Track

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 |
|---|---|---|
| **Inventory coverage rate** | % of known AI assets with a complete entry — owner, risk tier, and data mapped — against total assets identified across all four categories. | *Rising toward 100%* |

| Metric | What it measures | Target direction |
|---|---|---|
| **Mean time to discovery** | Average time between a new AI asset or agent going live and its appearance in the inventory. The clearest signal of whether discovery is continuous or reactive. | *Falling toward days, not months* |
| **Reconciliation gap rate** | % mismatch each month between automated tooling findings and the manual inventory. Directly measures whether the reconciliation step is actually happening. | *Falling quarter over quarter* |
| **Ownerless asset count** | Raw count of inventoried assets with no named business or technical owner. An absolute number, not a percentage — it reads as a to-do list, not an abstraction. | *Falling toward zero* |
| **Category coverage balance** | % breakdown of inventoried assets across the four categories. Flags when the inventory is really just a list of sanctioned platforms wearing a bigger label. | *Moving toward all four represented* |

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 inventory entry looks like, the 30/90/365-day build sequence, and the failure modes that derail this work most often.

## The Visibility Gap

Before building the inventory, your team needs to internalize why the gap exists. It isn't one blind spot — it's four distinct categories, each invisible for a different reason.

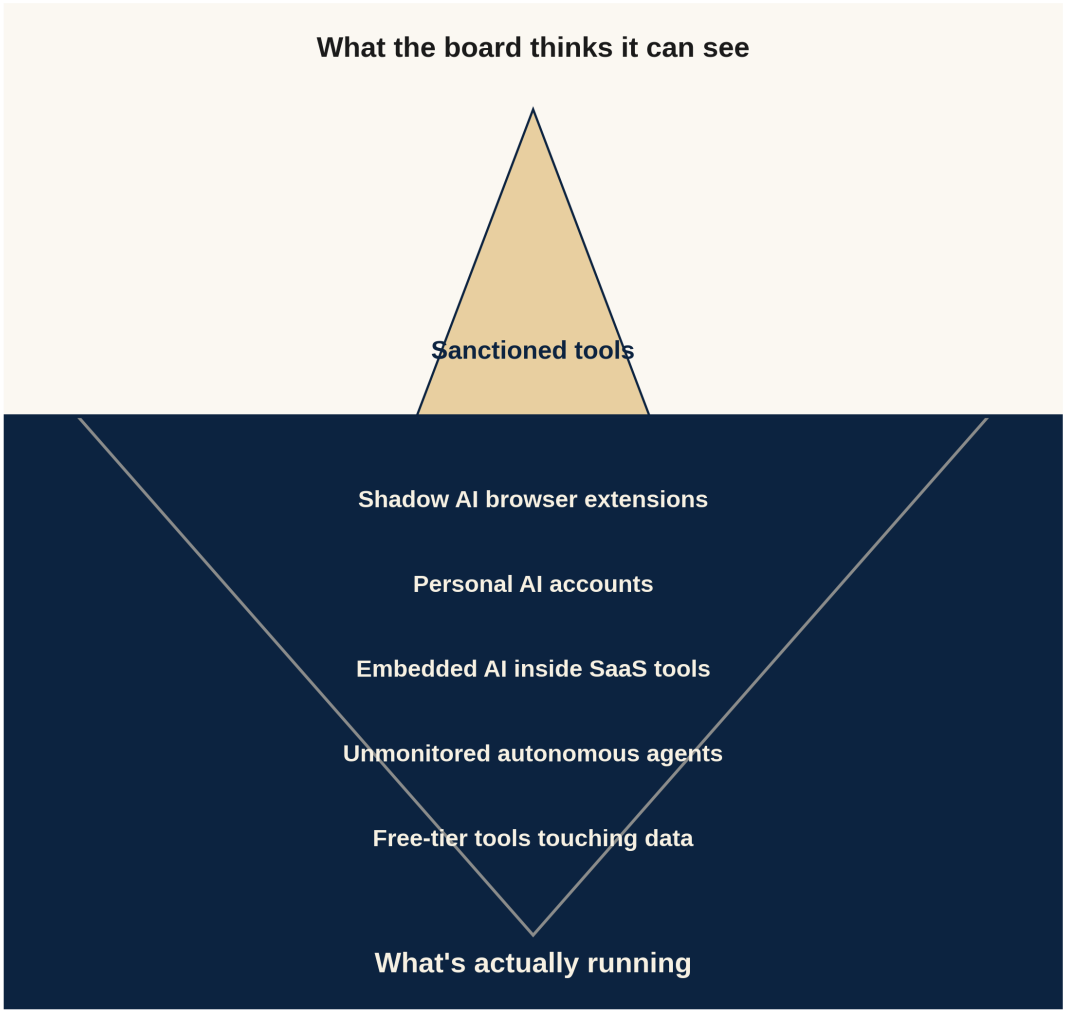


## The Four Asset Categories

### 1. Sanctioned platforms

Approved vendors and licensed AI tools that went through procurement. The easiest category to inventory — and the one teams mistake for the whole picture.

#### 2. Shadow / unsanctioned tools

Personal ChatGPT and Claude accounts, unapproved browser extensions, free-tier tools processing company data. Found through browser telemetry, DLP alerts, and expense report review — not through asking.

#### 3. Autonomous agents

Systems that take action, not just generate output. An agent that can take actions in your environment needs to be in your inventory with the same rigor as a human user — documented permissions, a named owner, a regular access review, and a decommissioning process.

#### 4. Embedded AI features

The most invisible category, and often the most data-sensitive. AI features enabled — sometimes by default — inside tools you already trust: summarization in your SIEM, transcription in your video platform. The security review happened before these features existed. Capturing this category requires vendor contract review and product changelog monitoring, not just an asset scan.

## 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 expense report that found what the scanner didn't

A CISO at a mid-size financial services firm had already run a CASB-based AI discovery sweep and reported “complete visibility” to the board. Three months later, finance flagged a routine expense reimbursement for an AI writing tool — paid on a personal card, expensed as “software,” used by a team drafting client-facing investment summaries. The tool had never touched the corporate network the CASB was watching. The fix wasn't better tooling; it was adding finance to the quarterly inventory review, a stakeholder group security had never looped in before.

#### The agent nobody remembered building

A healthcare CISO's team discovered, during an unrelated access review, a claims-processing automation script with standing write access to a production database — built two years earlier by a data science intern, never decommissioned, never assigned an owner after the intern's contract ended. It had been running unmonitored the entire time. Nothing had gone wrong. That was luck, not governance — and it became the case study the CISO used to get budget approved for an agent inventory and access review process the following quarter.

#### The vendor update that changed the data flow

An insurance CISO's organization had approved a video conferencing platform years earlier through a standard procurement review. A routine contract renewal call surfaced that the vendor had quietly added AI-powered meeting transcription and summarization — enabled by default — months prior. Sensitive claims discussions had been processed and stored by a third-party AI feature with no new security review, no updated data processing agreement, and no one in security aware it existed. The chapter's lesson held: the original approval predates the risk.

> ***The pattern across all three: the gap wasn't a missing tool. It was a missing stakeholder, a missing owner, or a missing trigger for re-review. Fix the process gap, and the tooling gap gets much smaller.***

## Your Inventory Framework

For every asset, capture identity, ownership, data, risk, technical, and control fields. The goal is not complete entries on day one — it's a structure you build out over time, highest-risk assets first.

## 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 that no tool can infer. Automated methods surface scale and blind spots no interview will uncover. A real inventory uses both, run in parallel, reconciled against each other on a set cadence.

### Manual methods — run these first

- Stakeholder interviews: 30-minute conversations with procurement, finance (expense report review catches AI subscriptions paid on personal cards), individual business unit leads, and your vendor management team. This is where you learn about tools that never touched a security review.
- 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 — most product updates never trigger a re-review.
- Business unit self-attestation survey: a short, recurring form (quarterly) asking each team to list AI tools in active use. Low completion rates are themselves a signal — that's the highest-risk group to interview directly.
- Expense and procurement report sweep: search for recurring charges to known AI vendors (OpenAI, Anthropic, Midjourney, and similar) on corporate cards — this catches paid shadow AI that free-tier discovery tools miss.
- Structured changelog review: assign someone to read release notes for your top 20 SaaS vendors monthly. AI features ship quietly, and this is often the only place they're disclosed before a user tells you.

### Automated methods — run these continuously

- Browser telemetry / DLP: your existing DLP or endpoint agent can usually be tuned to flag traffic to known AI domains (chat.openai.com, claude.ai, and similar) — this is the fastest way to surface unsanctioned personal-account use.
- CASB / SSPM platforms: cloud access security brokers and SaaS security posture management tools discover connected and OAuth-authorized apps automatically, including AI tools employees have granted access to company data.
- API and network egress monitoring: flag outbound calls to LLM provider APIs directly from your network or cloud environment — this catches engineering-built integrations that never went through procurement.
- Data-layer discovery platforms: map where sensitive data lives and what's reading from it — including AI tools and agents with access, even when the tool itself wasn't separately inventoried. This is the automated method most likely to catch what the others miss, because it starts from data access rather than the tool's name.
- 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 inventory has logged. Every unmatched entry on either side is a gap — either a tool no one interviewed about, or a tool your scanning missed. That reconciliation list is more valuable than either source alone.***

## Implementation Guide — 30 / 90 / 365

### The 30-day foundation

Week 1 is stakeholder mapping — identifying who actually has visibility into AI deployment across your organization (a different group than your usual IT contacts: procurement, individual business unit leads, and your vendor management team). Weeks 2–4 focus on capturing your highest-risk exposure: known sanctioned platforms plus any shadow AI already flagged by DLP or browser telemetry. The goal isn't completeness — it's a credible first version and a repeatable capture process.

### The 90-day build-out

Extend capture across all four categories. Run a structured discovery sweep for embedded AI features via vendor contract review and changelog monitoring. Assign named business and technical owners to every entry — no orphaned assets. Establish the review cadence you'll hold going forward.

### The 12-month target

Discovery becomes continuous, not episodic — new AI deployments trigger inventory entries automatically as part of procurement and engineering workflows. The inventory becomes a standing board-reportable artifact, reviewed on a fixed schedule, with open remediation items tracked to closure.

## Common Failure Modes

- Treating this as a one-time project instead of a maintained process — the inventory is stale within a quarter without a review cadence.
- Scoping only sanctioned platforms and reporting that as “complete” — this is the gap that surfaces worst in an audit or incident.
- No named owner per asset — an inventory entry without an accountable owner doesn't get maintained or remediated.
- Building the tracker before mapping stakeholders — teams that skip Week 1 spend months rebuilding trust with the business units who actually hold the visibility.

## This Week's Action

> ***Identify the three people in your organization outside of IT/Security who would know about an AI deployment before you would. Set up 30 minutes with each of them this week. That conversation is where your real inventory starts.***

## Chapter Toolkit

This chapter includes editable Board roadmap slides and the AI Asset Inventory Tracker (.xlsx) — a ready-to-use workbook with dropdown validation, example entries, and an auto-calculating summary dashboard by category and risk tier. Provided alongside this chapter.

SUPPORTED BY

**VARONIS**

**A note from Moriah**

*I've spent three CISO roles trying to answer the question this chapter opens with, and nobody solves it with a spreadsheet alone. Varonis Atlas is the clearest version of the fix I've seen: it continuously discovers AI assets — including shadow AI — and builds a lineage graph tying every model, agent, and copilot back to the sensitive data it can actually reach. It layers AI-SPM scanning, runtime guardrails, and third-party AI risk assessment — matching each vendor's AI Bill of Materials against real access — into that same inventory. That's why it catches what other methods miss: it starts from data access, not the tool's name. Varonis is supporting this chapter for that reason — and Chapter 3, where we go deep on data security, for the rest of the story.*