

GOVERNANCE PLAYBOOK, 2027

The ten questions you are about to be asked about internal AI, and how to produce each answer.

Written for the technology leader who approved the AI and now owns the number. Finance, HR, Legal and the CEO each hold one piece of the question, and all four arrive at your desk for the answer. This playbook names who asks each question, what they actually need, what usually gets sent back, and how your team produces an answer that holds up the second time it is asked.

WRITTEN FOR

The CIO, CTO, VP Engineering or head of IT at a company with a few hundred people or more, where staff already use AI tools, approved or not, and where the software the company already pays for has switched AI features on.

ALSO FOR

The CISO, the IT operations or FinOps lead who will build the record, and the finance, HR and legal leaders who will ask for it. Each can read the questions to know what to ask technology.

NOT FOR

AI inside a product you sell. That is the AI-native CTO playbook. The full legal register lives in the AI compliance playbook for counsel. This document describes what published sources say. It is not legal advice.

The question comes from four directions. The answer lives on your estate.

Finance wants the spend and who approved it. HR wants to know which tools people actually use, and which ones touch decisions about people. Legal wants to know what could be reconstructed after an incident. The CEO wants to know what any of it returned. None of them can produce the answer. Your team can, and most technology teams cannot yet either.

35%

of organizations report fully visible and actively monitored AI operating costs. 13% discover the cost only after billing.

KPMG, Global AI Pulse Q2 2026, 24 June 2026

52%

of organizations say there is no clear owner of AI cost. Only 13% have even basic visibility into actual AI spending.

Harness, State of AI in FinOps, 29 July 2026

78%

of IT leaders were hit with unexpected charges from consumption-based or AI pricing. 61% cut projects because of unplanned cost increases.

Zylo, SaaS Management Index, 29 January 2026

16%

of security leaders govern AI access to core systems effectively, while 71% say AI tools already reach systems such as Salesforce and SAP.

Saviynt and Cybersecurity Insiders, 21 April 2026

Read the first card twice. In the same KPMG survey, organizations with full cost visibility report established AI return at 15%, against 3% for those without it. The record is not overhead. It is the difference between internal AI that gets funded next year and internal AI that gets cut.

There are two pools of internal AI use, and only one of them needs to be locked down.

POOL ONE

Use that should run loose

Pilots, experiments, a team trying a tool, an engineer testing a model against a sample. Nothing in this pool has touched client material, an employee record or a deadline. The variance is healthy, and clamping down on it teaches people to route around you.

Know roughly what it costs and who is in it. That is enough.

POOL TWO

Use that should be locked down

AI that a team runs to do the actual work, on client material, on employee records, in a decision about a person, or inside a workflow a deadline depends on. From that day, the tool needs an approver, a named owner, a cost attributed by team and application, and a figure that ties to the invoice.

This is the pool the other four desks are asking about.

Almost every organization reports one number for both pools. That is how a pilot quietly becomes a dependency nobody approved, and how a policy can be signed by everyone and enforced on no one.

So the useful question is never what is our AI spend. It is which pool is this in, who approved it, and does the number tie to the bill.

WHAT HR AND LEGAL ARE WORKING AGAINST

Several of the questions now have a statute behind them, and the record they need is yours.

You do not need to interpret these. Counsel does. You need to know which ones turn on a record your team produces, and how long that record has to last. The entries below match the register in the companion AI compliance playbook for counsel, checked on the dates it states.

California FEHA automated-decision system regulations

In force 1 October
2025

An employer's use of an automated-decision system in employment decisions is subject to the Fair Employment and Housing Act. Anti-bias testing is relevant evidence. The inputs, outputs, criteria and testing results are kept for at least four years, and the employer can be liable for a software vendor acting on its behalf. For you, that is a retention period on the record for every tool that screens, ranks or scores people.

McDermott Will and Emery, 26 September 2025. Mayer Brown, 2 September 2025

Illinois HB 3773

In force 1 January
2026

Prohibits employer use of AI that has the effect of discriminating on a protected class, and requires notice before AI is used in recruitment, hiring, promotion, discharge, discipline and the other listed purposes. For you, HR cannot give the notice without knowing which tools touch those decisions.

Crowell and Moring, September 2024. National Law Review, May 2025

Connecticut Public Act 26-15

In force 1 October 2026. Notice from 1 October 2027

Since 1 October 2026, using automated employment decision technology is not a defense to a discrimination claim, and a WARN notice must disclose whether layoffs relate to the employer's use of AI. From 1 October 2027, employers give written notice naming the purpose, the decisions affected, the tool's trade name and the categories of personal data it processes. For you, the trade name and data categories are a register row.

Epstein Becker Green, 8 June 2026. Proskauer, 4 June 2026

Colorado SB 26-189

Employer duties from 1 January 2027

Deployers whose automated decision-making technology materially influences consequential decisions, including employment, owe notice, a plain-language explanation within 30 days of an adverse decision, human review where commercially reasonable, and records kept at least three years after the decision.

Proskauer, 1 June 2026. Constangy, 30 July 2026

SEC Form 8-K Item 1.05

Adopted 26 July 2023

A material cybersecurity incident is disclosed within four business days of the materiality determination. On 11 May 2026 an employee's unauthorized AI use triggered the first such filing. For you, a record that has to be assembled after the incident cannot be assembled in four days.

Wilson Sonsini, 28 May 2026

Each one names who asks it, what a defensible answer looks like, and how to produce it.

The middle box is what usually gets sent back today. None of it is dishonest. It is what a competent person sends when the instrumentation cannot do better, and it is the reason the same question comes back next quarter.

QUESTION 01 ASKED BY FINANCE

What did AI cost us last month, in dollars, including the AI inside the software we already pay for?

The baseline, and it is wider than the API bill. Microsoft raised Microsoft 365 commercial prices by 5% to 43% on 1 July 2026 and folded Copilot Chat into base plans, so AI cost now arrives as a list-price increase rather than an add-on (MarketScale, July 2026). GitHub moved every Copilot plan to usage-based billing on 1 June 2026. The subscription line is where most of the surprise now lives.

A DEFENSIBLE ANSWER

One figure, broken out by provider and by application, that includes every subscription with AI features switched on and every expensed tool.

A PLAUSIBLE ANSWER

The API provider invoices. Accurate as far as they go, and they miss the enterprise applications that switched AI on and the tools paid for on cards.

HOW TO PRODUCE IT

One inventory: every provider account, every subscription with AI enabled, every expensed tool. One owner for the list. Refreshed monthly, not assembled on request.

QUESTION 02 ASKED BY FINANCE AND THE AUDITOR

Does that number tie to the invoices we actually paid?

The most important question in the playbook. Usage logs produce a computed figure, and the invoice is what was paid. A computed number has no way of knowing it is wrong, and the gap between the two is where untracked usage and price drift hide. 13% of organizations discover their AI cost only after billing (KPMG, June 2026).

A DEFENSIBLE ANSWER

A tie-out on a stated cadence, and the last time the two numbers did not match, with the reason.

A PLAUSIBLE ANSWER

A dashboard screenshot. A dashboard is a calculation. It is not a reconciliation, and it cannot tell you when it is wrong.

HOW TO PRODUCE IT

Monthly, line by line, attributed spend against each provider invoice, run with Finance. Every gap named as one of two things: usage you were not tracking, or a price that moved.

QUESTION 03 ASKED BY HR

Which AI tools are our people actually using, and how do you know?

Not the approved list. The measured one. Two thirds of office professionals have used an AI tool at work they believed violated policy (PagerDuty, June 2026). 67% of AI access from corporate devices runs through non-corporate accounts, and the most common data submitted to external models was company source code (Verizon DBIR, May 2026).

A DEFENSIBLE ANSWER

A measured list of tools in use, including personal accounts on corporate devices, with the method behind it and a date. A small honest number beats a confident zero.

A PLAUSIBLE ANSWER

The approved list, or the policy. Both are statements of intent. Neither is a reading.

HOW TO PRODUCE IT

Sign-on logs, expense lines, network and DNS telemetry, and the provider accounts themselves. Count what appears, compare it to the approved list, and report the difference, not the list.

QUESTION 04 ASKED BY HR AND LEGAL

Who approved each tool, for what data, and is that written down?

The approval trail. When a customer's security review, an auditor or a regulator asks, the first thing wanted is not the spend. It is who decided this tool could touch this class of data, and when. It now carries a date as well. From 1 October 2027, Connecticut employers must tell people the trade name of each tool used in employment decisions and the categories of personal data it processes. Most organizations have a procurement ticket. Very few have a register.

A DEFENSIBLE ANSWER

A register per application: owner, approver, date, the data classes it may process, and whether it makes or shapes decisions about people. Category recorded, never content.

A PLAUSIBLE ANSWER

"IT approved it," or a link to the purchase order. Neither says what the tool was approved to do.

HOW TO PRODUCE IT

One row per application in the inventory from question one. Fill the approver, data-class and people-decision columns for the production pool first. Leave the pilot pool with an owner and a date.

QUESTION 05 ASKED BY FINANCE

What does each team cost, and does that figure tie to the invoice?

The chargeback question, and the one that turns a total into an instrument. Only 13% of organizations have even basic visibility into actual AI spending, and organizations estimate 26% of it is wasted (Harness, July 2026). You cannot find the waste in a total. You find it in the team and the application that produced it.

A DEFENSIBLE ANSWER

Cost by team and by application, tied to the invoice, with the top three and the bottom three named.

A PLAUSIBLE ANSWER

An allocation by headcount or by licence count. It sums correctly and says nothing about which team drove the cost.

HOW TO PRODUCE IT

Attribute at the request level to the team and application that generated it, then tie the sum to the bill. Anything that cannot be attributed goes on its own line, named, so the gap is visible instead of spread.

QUESTION 06 ASKED BY THE CEO

Which of this is a pilot, and which has quietly become something the business depends on?

The two pools, asked out loud. Nothing marks the transition. A team tries a tool in March, by June a client deliverable depends on it, and the P&L never noticed. That is how spend nobody approved becomes spend nobody can stop. 49% of organizations scaled back, narrowed, delayed or paused AI agent deployments after costs outran value (KPMG, June 2026).

A DEFENSIBLE ANSWER

Two lists, and the written rule that moves an application from one to the other, with the name of whoever applies it.

A PLAUSIBLE ANSWER

"We have several pilots running."
That is one list with no exit condition.

HOW TO PRODUCE IT

Write the rule once. The day a workflow touches client material, employee records, a decision about a person or a deadline, it moves to production and gets an approver, an owner, a budget and an alert. Then apply it to the inventory.

QUESTION 07 ASKED BY FINANCE AND THE CEO

If a team's usage doubled tomorrow, would anyone know before the invoice?

72% of organizations had an unexpected AI cost spike in the past year, 33% more than once, and only 20% could identify the cause of a spend doubling within hours (Harness, July 2026). 40% of organizations in a separate survey had an unexpected AI cost that needed board-level escalation (Mavvrik with Benchmarkit, July 2026). Uber's engineering organization went from 32% to 84% adoption of one coding tool in three months and used its entire annual AI budget in four (Fortune, via Forbes, July 2026). The spike is not the failure. Finding out from the invoice is.

A DEFENSIBLE ANSWER

A budget per team and per application, a threshold, an alert, and the name of the person who receives it the same day.

A PLAUSIBLE ANSWER

"We review spend monthly." A monthly review finds the spike after it has been paid for.

HOW TO PRODUCE IT

One threshold per production application, even a crude one, routed to the owner in the register. Record the largest surprise so far and how it was found. That record is the answer next time.

QUESTION 08 ASKED BY LEGAL

If an employee took a harmful action through an AI tool tomorrow, what could we reconstruct?

This stopped being a hygiene question in 2026. 43% of security incidents now involve unsanctioned AI, more than double the prior year, and over 66% of breached organizations had no governance process to limit it (IBM, July 2026). In May 2026 a bank filed the first SEC disclosure triggered by an employee's unauthorized AI use, after customer names, Social Security numbers and dates of birth went through an unapproved app. The rule allows four business days.

A DEFENSIBLE ANSWER

Who used which application, on which team's behalf, at what scale, on which dates, with retention agreed in writing. The category of data, with access to any content governed by a written rule.

A PLAUSIBLE ANSWER

The policy and the training completion log. Those describe what people were told. They reconstruct nothing.

HOW TO PRODUCE IT

Keep the attribution record from question five running continuously, with retention set by Legal. For tools in employment decisions that means at least four years in California and three in Colorado. It has to exist before the incident.

QUESTION 09 ASKED BY FINANCE AND PROCUREMENT

What are we exposed to if a provider reprices, or moves us from seats to usage?

Concentration risk, and it is no longer hypothetical. 79% of IT leaders met a SaaS price increase at renewal in the past twelve months (Zylo, June 2026). Two of the most common AI line items changed pricing model within a month of each other in 2026. Microsoft 365 folded AI into base plans, and GitHub Copilot moved every plan to usage billing.

A DEFENSIBLE ANSWER

Share of spend by provider, the renewal date for each, which contracts moved from per-seat to usage, and what a material increase from the largest one does to the budget.

A PLAUSIBLE ANSWER

"We could switch." Possibly true, and silent on how long it takes and what breaks.

HOW TO PRODUCE IT

One table from the inventory: provider, share of spend, renewal date, pricing model, and the run number for a stated increase. Flag every subscription whose pricing changed in the last year.

QUESTION 10 ASKED BY THE CEO AND THE BOARD

What did the internal AI return, and who owns that number?

The closing question, and the one the funding decision runs on. 87% of finance leaders feel pressure to tie AI spend to outcomes within a year, and only 22% have done so (CloudZero, via CFO Dive, July 2026). Adoption figures do not answer this. Seats and logins describe procurement, not return.

A DEFENSIBLE ANSWER

Cost per team and per workflow, set against a stated outcome for the workflows the business depends on, and one name who owns the number.

A PLAUSIBLE ANSWER

Adoption, usage minutes, and a survey of how people feel about the tools. Useful signals. None of them is a return.

HOW TO PRODUCE IT

Pick the three production workflows that matter most and name the outcome metric for each. Put the tied-out cost against it every month, in a standing report with your name on it. Forward it before you are asked.

BUILDING THE RECORD

Six moves, in this order. Each one makes the next one possible.

All ten questions draw on the same record. Build it once and every desk gets its answer from the same place. Build it per request and you will be assembling it again next quarter.

01 Inventory

Every provider account, every subscription with AI switched on, every expensed tool, every personal account found on a corporate device. One list, one owner, refreshed monthly. Answers questions one and three.

02 Register

For each application, the owner, approver, date, data classes allowed, which pool it sits in, whether it makes or shapes decisions about people, and the legal entity and country it is used in. Category recorded, never content. Answers questions four and six.

03 Attribute

Every request tied to the team and application that generated it. Unattributed usage gets its own line so the gap stays visible. Retention agreed with Legal in writing, at least as long as the longest rule that applies. Answers questions five and eight.

04 Tie out

Monthly, Finance compares the attributed record to each provider invoice, line by line, and names every gap as usage nobody was tracking or a price that moved. This is the move that turns a computed number into a checked one. Answers question two.

05 Alert

A budget and a threshold per production application, routed to the owner in the register. Record every surprise, its cause, and how it was found. Answers questions seven and nine.

06 Report

The same five numbers every month: cost by team, cost by application, the gap to invoice, unapproved usage found, and production spend against the stated outcome. Forwarded to Finance, HR and Legal before they ask. Answers question ten.

READING THE ANSWERS

**Count the ones you could answer this afternoon,
not the confident ones.**

Score a question as answered only if the figure, the name or the cadence exists today, with a document or a date attached, and you could produce it without a week of notice.

EIGHT TO TEN

Ahead of nearly everyone.

The remaining gap is usually question two, because the tie-out needs an attributed record and someone in Finance who runs it every month. Close it and the record stands up to an auditor.

Send the monthly report before anyone asks for it. That is the whole job now.

FOUR TO SEVEN

Normal, and worth a decision.

Do not try to close every gap at once. Build the inventory and the register first, because every other answer draws on them. Then tie out once, by hand if you have to.

Put a date against questions one, two and four, and tell Finance, HR and Legal what it is.

THREE OR FEWER

An instrumentation problem, not a competence one.

Only 35% of organizations report fully visible and monitored AI costs, so this is where most of the market sits. The risk is not a weak team. It is an outside party finding the number first.

Start with the inventory. It changes the answer to three of the ten in an afternoon.

One caution. "We do not have that yet, and here is when we will" is a stronger answer than a confident number nobody has tied to an invoice. The second one is the one that becomes a surprise later, with your name on it.

TAKE THIS INTO THE ROOM

Internal AI. Ten questions for the technology leader.

Print this, or send it to Finance, HR and Legal before the meeting so everyone is asking from the same sheet.

- | | | |
|-------|--|----------------------|
| 01 | What did AI cost us last month, including the AI inside the software we already pay for? | FINANCE |
| <hr/> | | |
| 02 | Does that number tie to the invoices we actually paid? | FINANCE, AUDIT |
| <hr/> | | |
| 03 | Which AI tools are our people actually using, and how do you know? | HR |
| <hr/> | | |
| 04 | Who approved each tool, for what data, and is that written down? | HR, LEGAL |
| <hr/> | | |
| 05 | What does each team cost, and does that figure tie to the invoice? | FINANCE |
| <hr/> | | |
| 06 | Which of this is a pilot, and which has quietly become something the business depends on? | CEO |
| <hr/> | | |
| 07 | If a team's usage doubled tomorrow, would anyone know before the invoice? | FINANCE, CEO |
| <hr/> | | |
| 08 | If an employee took a harmful action through an AI tool tomorrow, what could we reconstruct? | LEGAL |
| <hr/> | | |
| 09 | What are we exposed to if a provider reprices, or moves us from seats to usage? | FINANCE, PROCUREMENT |
| <hr/> | | |
| 10 | What did the internal AI return, and who owns that number? | CEO, BOARD |
| <hr/> | | |

Score a question as answered only if the figure, the name or the cadence exists today, with a document or a date attached.

WHERE THIS CAME FROM

We build the record the other four desks ask for.

Crystil is an independent third party, SOC 2 compliant. One line of code prices every model call by tokens and model, and attributes the cost to the team, application or customer behind it.

Workflows group many model calls into one task, so a whole outcome carries one cost. Pulse routes a team's AI coding tools, Claude Code, Codex and Gemini CLI, through Crystil with no SDK.

Sentinel is opt-in and off by default. When a customer switches it on, it checks prompts against the rules the customer sets before they reach the model, and can stop a prompt carrying information the company has ruled out, such as employee records or privileged material, and notify HR and Legal.

Crystil does not reconcile your invoices. It provides the attributed record your finance team reconciles against them. The inventory, the register and the return metric stay your team's work.

[Talk to an AI measurement expert](#)

crystil.com

COMPANION GUIDES

The AI compliance playbook for counsel asks these questions from the legal side, with the full register. **The AI playbook for HR leaders** covers the tools that shape decisions about people and the notices they need.

Sources

1. KPMG, Global AI Quarterly Pulse Survey Q2 2026, 24 June 2026.
2. Harness, 2026 State of AI in FinOps, 29 July 2026. Vendor research, 700 leaders.
3. Zylo, 2026 SaaS Management Index, 29 January 2026, and SaaS Pricing Trends, 9 June 2026. Vendor research.
4. Saviynt and Cybersecurity Insiders, 2026 CISO AI Risk Report, 21 April 2026. Vendor research.
5. MarketScale, 15 July 2026, on Microsoft 365 pricing. GitHub, Copilot usage-based billing, 27 April 2026.
6. PagerDuty, shadow AI survey, 11 June 2026. Verizon 2026 DBIR, via National Law Review, 28 May 2026.
7. Mavrik with Benchmarkit, 2026 State of AI Cost Governance Report, 29 July 2026. Vendor research.
8. Fortune, via Forbes, As Token Costs Plunge, Enterprise AI Providers Face A New Margin Squeeze, 28 July 2026.
9. IBM, Cost of a Data Breach Report 2026, via Cybersecurity Dive, 29 July 2026.
10. Wilson Sonsini, WSGR Data Advisor, 28 May 2026. SEC press release 2023-139, 26 July 2023.
11. CloudZero, via CFO Dive, 22 July 2026. Vendor research.
12. California FEHA ADS rules: McDermott Will and Emery, 26 September 2025, Mayer Brown, 2 September 2025. Illinois HB 3773: Crowell and Moring, September 2024, National Law Review, May 2025.
13. Connecticut PA 26-15: Epstein Becker Green, 8 June 2026, Proskauer, 4 June 2026. Colorado SB 26-189: Proskauer, 1 June 2026, Constangy, 30 July 2026.

Crystil. v1.0, 7 October 2026. Every figure is published and dated. Nothing here is modelled, estimated or invented. Legal entries follow the counsel playbook register. This document describes what published sources say. It is not legal, tax or investment advice.