
AI GOVERNANCE PLAYBOOK

The AI Governance Checklist

Seven principles and a five-step sequence for teams scaling with AI in their workflows.



You Don't Need Less AI. You Need To Be Able To See It.

Every growing engineering team ends up running AI somewhere in its workflow, whether anyone signed off on it or not. The real question isn't whether your team uses AI. It's whether you can still say, with confidence, who did what, which tools touched your data, and what it's actually costing you.

That gets harder exactly when your team is growing: more people, more time zones, more tools nobody centrally approved. Governance that worked for a ten-person team doesn't hold up on its own once you're fifty.

This playbook sets out seven principles worth building into how your team operates, and a five-step sequence for putting them into practice. None of it requires replacing your current stack. It requires deciding, before growth forces the decision for you, what a well-governed team does by default.

HOW TO USE THIS

Page 3 is a one-page version you can pin up or check off directly. Pages 4 to 7 walk through each principle, the five-step sequence, and where to start.

-
- | | |
|----|-------------------------------|
| 03 | Seven principles, at a glance |
|----|-------------------------------|
-
- | | |
|-------|----------------------------------|
| 04-05 | What this looks like in practice |
|-------|----------------------------------|
-
- | | |
|----|----------------------|
| 06 | A five-step sequence |
|----|----------------------|
-
- | | |
|----|----------------|
| 07 | Where to start |
|----|----------------|

Seven Principles, At a Glance

Tick these off as you build them in. Full explanations start on the next page.

- ☐ Approve an AI stack, and make it the easy path 01
- ☐ Fold AI activity into normal management visibility 02
- ☐ Connect usage, cost, and risk in one view 03
- ☐ Separate what a person built from what AI generated 04
- ☐ Build audit trails in by default 05
- ☐ Control data before it reaches the model 06
- ☐ Give autonomous agents runtime guardrails 07

Print this page. Everything after it is the detail behind each line.

What This Looks Like In Practice

01

Approve an AI stack, and make it the easy path

Decide which assistants, coding agents, automation tools, and model providers your team is approved to use. Then make sure that approved path is genuinely easier to reach than an unapproved one. People default to whatever's fastest to open, so the fastest option has to be the one you can actually see.

02

Fold AI activity into normal management visibility

Don't treat AI as a separate category to review later. Track how it shows up in delivery, code review, incident response, and anything customer-facing, right alongside everything else your team already reports on.

03

Connect usage, cost, and risk in one view

A dashboard that only shows who's using what isn't enough on its own. Tie usage to spend, data exposure, security signals, and review patterns, so a spike in one shows up against the others instead of sitting in its own silo.

04

Separate what a person built from what AI generated

Record what a person wrote, what a model generated, what got reviewed, and what actually shipped. Building this distinction in from day one is far easier than trying to reconstruct it after the fact.

PRINCIPLES 5-7 CONTINUE OVERLEAF

What This Looks Like In Practice

05

Build audit trails in by default

For every AI-assisted workflow, log who ran it, what they asked for, which system it touched, which model or agent handled it, what came out, and what happened next. This is what turns a written policy into something you can actually demonstrate.

06

Control data before it reaches the model

Sensitive data should be caught and governed before it's pasted into a prompt, uploaded to a tool, or read by an agent, not after.

07

Give autonomous agents runtime guardrails

Scoped identities, limited permissions, logging, and points where a human has to approve something all limit how much damage an agent can do if it goes off track. Think of a new agent like a new hire on day one: it gets what it needs to do the job, not every key in the building.

These seven hold up whatever tools you're running today, and whatever you switch to next.

A Five-Step Sequence

Whether you start with people, with AI, or with both at once, this order tends to hold up.

1

Take inventory

Find every AI tool and agent already in use across your team: official apps, unofficial ones, browser extensions, and any custom agent workflows nobody formally signed off on.

2

Classify the risk

Sort AI activity by what it touches: public information, internal business data, customer data, source code, credentials, regulated data, and anything with access to production systems.

3

Set policy boundaries

Decide what people and agents are allowed to access, what data can leave your systems, when a human has to review something before it ships, and which actions an agent can take entirely on its own.

4

Build in visibility

Bring usage metrics, agent activity logs, data-loss events, cost tracking, and workforce activity into one place that actually shows what's happening across your team, not five dashboards that don't talk to each other.

5

Review and improve

Use what you're seeing to narrow your approved tools, fix workflows that aren't working, cut AI spend that isn't earning its keep, catch risky patterns early, and coach your team as you go.



WHERE TO START

Most of this you can start today, with the team you already have.

If you'd rather have a second pair of eyes on where things actually stand,
that's a conversation we're glad to have.

Work With Us

codeclan.com.au/contact

