



AI Governance, Security, and Compliance.

The three disciplines an organization needs in place before implementing AI and how they fit together.

Governance decides what you allow and who is accountable for it. Security protects the systems and the data behind it. Compliance is what lets you see what's actually happening. Some organizations have these disciplines in place with no clear line between them, others have a piece of each.

WHY THIS COMES FIRST

AI usually arrives before anyone approves it.

Most organizations do not decide to adopt AI. Someone turns on a feature inside software the company already pays for; a team starts using a free tool for customer email, a vendor ships an assistant into a product the business depends on. By the time leadership asks who approved it, the honest answer is nobody, and the second is that nobody knows the full list.

That is not a technology failure. It is a management gap, and it closes in three moves: decide what is allowed and who owns it, protect the systems and data involved, and keep records that let you answer a hard question quickly. Those three moves are **governance**, **security**, and **compliance**. They are not substitutes for one another, and skipping any one of them shows up later as an incident you cannot explain, or an opportunity you had to decline.

What each one produces, and what its absence costs:

	Governance Should we do this, under what rules, and who is accountable?	Security How do we prevent, detect, contain, and recover from harm?	Compliance What obligations apply, and can we show we met them?
WITH IT	Strategy, named owners, risk appetite, approved and prohibited uses, decision rights, an oversight cadence.	Access controls, tool permissions, data boundaries, monitoring, incident response, safe fallback.	Obligation register, controls mapped to requirements, assessments, retained records, audit evidence.
WITHOUT IT	Uncoordinated adoption, nobody accountable, and tools nobody approved doing work nobody checked.	Data leaving quietly, outputs that can be manipulated, actions you did not authorize, operations you cannot restore.	Regulatory, contractual, insurance, and procurement exposure you cannot answer on demand.

RISK TIERING

Not every AI use needs the same scrutiny.

The fastest way to make governance feel unbearable is to apply the same process to a marketing draft and a hiring decision. Sorting by consequence first lets most AI use move quickly, so the small number of consequential systems get real attention.


Low	LOOKS LIKE Internal drafting and research on information that is not sensitive, with nothing sent outside the company and an output a person can check at a glance.	TREATMENT Log it in the register, keep it on an approved tool, cover it in general training, and revisit on the annual cycle. Days of effort, not weeks.
Moderate	LOOKS LIKE Customer-facing content, sensitive internal data, workflow recommendations, limited integrations.	TREATMENT Risk assessment, vendor and privacy review, testing before launch, monitoring by your IT provider, one named owner.
High	LOOKS LIKE Employment, healthcare, credit, legal, safety, eligibility, significant finance, or AI that acts on its own outside the company.	TREATMENT Formal impact assessment, legal review, independent testing, strict access, a human decision-maker, executive approval.
Prohibited	LOOKS LIKE Uses that conflict with law, rights, contract, policy, ethics, or your stated risk appetite.	TREATMENT Blocked at acquisition and deployment, monitored for shadow use, with the rationale written down.

Agentic AI and risk tier. A chatbot produces text and a person decides what to do with it. An agent retrieves documents, uses credentials, updates records, sends messages, and can start transactions. Once AI can act rather than only answer, the question stops being output quality and becomes what it is permitted to reach, what it can do without asking, and whether the change can be undone. Anything with that capability starts at moderate and usually lands at high.

ROLES AND RESPONSIBILITIES

Who does what, and when.

AI governance most often stalls because nobody agreed who owns which part, so the framework gets written and never implemented. Setting the division of labor at the start is what makes the rest hold.



Separation of duties. This is a standard control, not a stylistic preference. When one party defines a control, implements it, and then confirms that it works, no independent verification exists, and the evidence carries less weight. Keeping those three roles in separate hands is what makes verification work.