HERDING WOLVES

Five questions that tell you, in one conversation, whether a company has AI governance *or just AI experimentation.*

Most companies don't have an AI strategy. They have several, one per team, and each is reasonable on its own. **AI governance is how an organization decides how AI is allowed to be used.** Ask these five early in the first meeting. The answers show where the project is likely to break, and it's rarely the technology.

01 What AI tools are in use?

WHAT THE ANSWER TELLS YOU

Whether anyone can see the whole picture. Count in three places: tools people signed up for on their own, AI switched on inside software the company already approved, and AI built into whatever the company sells.

LISTEN FOR

A confident number from leadership. It's usually a guess. If people have been told "no," they've stopped telling anyone.

02 What data can they reach?

WHAT THE ANSWER TELLS YOU

The blast radius. A chatbot that answers questions and an agent connected to email, the CRM, and the file share are different risks, even from the same vendor. Guardrails define what a tool *won't* do. Access controls define what it *can't* do.

LISTEN FOR

"We use the enterprise version" as the whole answer, or nobody sure which accounts are personal and which belong to the company.

03 Where is human review required?

WHAT THE ANSWER TELLS YOU

Whether anything checks the work before it reaches a customer, the codebase, or the books, and whether that check was designed or just assumed.

LISTEN FOR

"Everywhere." That means nowhere in particular. Review that isn't designed wears away as people get comfortable.

04 Who approves new tools?

WHAT THE ANSWER TELLS YOU

Whether there's a path to yes. Without one, shadow AI is a certainty. A good answer names who decides, how fast, and when IT or security joins: at selection, before anything is signed.

LISTEN FOR

"Just use good judgment," or IT finding out after the contract is signed.

05 Who owns the outcomes?

WHAT THE ANSWER TELLS YOU

Who answers when an AI-made asset can't be protected, an agent changes the wrong record, or AI-written code ships a bug. One question for the broker: cyber policies were largely written for an outsider breaking in. When a tool the company *authorized* causes the loss, which policy pays: cyber, E&O, or neither?

LISTEN FOR

"The vendor." Or a pause, then a different name from each person in the room.

A company that can't answer these five is still experimenting with AI. **That's a fine place to start.** The gaps are the part of the project nobody scoped.

A SECOND SET OF EYES

Bring one company to a 30-minute conversation: yours, or a client you're implementing for. I'll walk these five questions against it with you and show you where it's likely to break.

Korey Mercier · Founder & Executive Advisor, Herding Wolves