AI Governance for SaaS: 5 Questions to Ask Before You Ship

Intro

You don't need a 50-page AI policy to start governing AI responsibly.

If you're a SaaS or AI-native team adding AI features, you can get a long way by asking five clear questions before you ship anything new.

These questions are designed for founders, product and engineering leads, and whoever owns security/privacy today.

Question 1: What problem is this AI feature actually solving—and for whom?

Before you think about models or prompts, get crisp on:

- Who this feature is for (end users, internal users, admins, etc.)
- What decision or workflow it's influencing
- How bad it is if it's wrong or misused (low, medium, high impact)

If you can't explain the use case and impact in a few sentences, you're not ready to ship.

Question 2: What data does it touch—inputs, context, training, and logs?

For each AI feature, be explicit about:

- Inputs: What users or systems send into the model
- Context: What extra data you attach (customer records, history, metadata)
- Training data: Whether you're training or fine-tuning on customer data
- Logs: What you store, for how long, and who can see it

You should be able to answer:

- Whose data is this?
- Is any of it personal, sensitive, or regulated (e.g., PHI)?

- Where does it live, and which vendors see it?

If you don't know, you can't make good security or privacy decisions.

Question 3: What do your model providers do with your data?

If you're using third-party models or APIs, read (really read):

- Their data-use and retention policies
- Whether they train on your prompts or outputs by default
- Their security and privacy commitments (and where they fall short)

You should know:

- Are they a vendor/subprocessor in your records?
- Do you have a DPA or equivalent terms in place?
- What happens if they change their terms or pricing?

If you wouldn't be comfortable explaining your provider choices to a big customer, you probably need a stronger story.

Question 4: How will users know they're using AI—and what are its limits?

Transparency is a core part of responsible AI:

- Make it clear when users are interacting with an AI-powered feature.
- Set expectations about accuracy, limitations, and when human review is required.
- Be honest about what you do with their inputs (training, product improvement, etc.).

Ask yourself:

- Would a reasonable user feel misled by how we present this feature?
- If something goes wrong, would we be comfortable pointing to our explanations and disclosures?

If the answer is "I'm not sure," tighten your messaging before you ship.

Question 5: Who owns this feature after launch?

AI features aren't "set and forget." You need clear ownership for:

- Monitoring metrics (errors, overrides, flagged outputs, complaints)
- Reviewing changes to models, prompts, or data sources
- Deciding when to roll back, retrain, or retire the feature

You should be able to name:

- A business owner (often product)
- A technical owner (engineering)
- A risk/compliance contact (security/privacy)

If ownership is fuzzy, issues will fall through the cracks.

Bringing It Together

If you can answer these five questions clearly, you're already doing more AI governance than many larger companies.

From there, you can:

- Formalize your answers into a simple AI risk assessment template
- Plug AI into your existing SDLC, risk management, and incident response processes
- Use the AI Feature Launch Checklist as a deeper pre-launch gate

You don't need a separate AI bureaucracy—you need clear thinking, a few good questions, and the discipline to ask them every time you ship something new.