

NexTrend.

RISK . CONTROLS . ADVISORY

AI Use Case Risk Intake *Form*

The 6 questions every company should answer before deploying any AI tool.

Grounded in the NIST AI Risk Management Framework

nextrendadvisory.com

WHY THIS EXISTS

Most companies don't govern their AI. They deploy it and hope.

AI gets adopted one tool at a time, usually with no one asking the basic risk questions until something goes wrong. This form is the gate that should sit in front of every AI deployment. It takes ten minutes, it forces accountability, and it creates a record you can defend later.

HOW TO USE IT

- 1 Run it before deployment.**
Any new AI tool, feature, or vendor model goes through these 6 questions first.
- 2 Name an owner.**
Every use case gets one accountable person. No owner, no approval.
- 3 Rate the risk.**
Use the matrix near the end to land on Low, Medium, or High.
- 4 Record the decision.**
Approve, add conditions, or reject. Write it down. Set a re-review date.

What is the use case, and who owns it?

WHY IT MATTERS

An AI use case with no named owner is an incident waiting to happen. Accountability is the foundation of governance: someone has to be answerable for this system before it touches anything real.

CAPTURE

- Use case name and the problem it solves
- Business owner (one accountable person, named)
- Requesting team and sponsor
- The specific AI tool, vendor, or model in scope

01

What data does it touch?

WHY IT MATTERS

Most AI risk is data risk. Sensitive data flowing into a model, especially a third-party one, is the fastest path to a breach or a privacy violation. You cannot govern what you have not classified.

CAPTURE

- Data going in and coming out (public, internal, confidential, PII, PHI, financial)
- Source of any training or reference data
- Where data is stored, and who can access it
- Privacy exposure under GDPR / CCPA

What happens if it is wrong?

WHY IT MATTERS

A wrong answer in a brainstorming tool is harmless. A wrong answer in a lending, hiring, or clinical decision is not. The worst realistic outcome, not the average one, drives the risk rating.

CAPTURE

- Worst realistic outcome of an incorrect or biased output
- Who is affected: customers, employees, the public
- Financial, legal, safety, and reputational impact
- Is the outcome reversible?

03

Where is the human in the loop?

WHY IT MATTERS

Human oversight is the single most effective control on AI risk. Autonomy without oversight turns a model error directly into a business decision, with no one positioned to catch it.

CAPTURE

- Is output reviewed before it acts?
- Who reviews, and do they have authority to override?
- Can the system take action on its own?
- What is the escalation path when output looks wrong?

Is it a third-party model, and what did you verify?

WHY IT MATTERS

You inherit the risk of every model you use but do not control. Vendor AI risk is third-party risk with a new shape, and the contract is where most of that risk is won or lost.

CAPTURE

- Vendor and model name
- Contract terms on data use and retention
- Security and compliance attestations (SOC 2, ISO 27001)
- Does your data train their model?

How will you monitor it over time?

WHY IT MATTERS

AI is not deploy-and-forget. Models drift, data shifts, and performance degrades quietly. Governance is continuous, not a one-time gate, so monitoring has to be owned before you go live.

CAPTURE

- Who owns ongoing monitoring
- Metrics for accuracy, bias, and drift
- Review cadence
- Incident response and a kill-switch plan

STEP 5

Rate the risk

Combine data sensitivity, impact if wrong, and autonomy. The highest of the three sets the floor. When in doubt, rate it up.

LOW

Public or internal data. Minor, reversible impact. A human reviews every output.

DECISION

Approve. Document and move on.

MEDIUM

Confidential data, or moderate impact, or only limited human review.

DECISION

Approve with conditions. Add controls and a review date.

HIGH

PII / PHI / financial data, or severe or irreversible impact, or autonomous action.

DECISION

Escalate. Formal review before any deployment.

STEP 6

Record the decision

This is the artifact that protects you later. Fill it in, store it, set the re-review date.

Use case name	
Business owner	
Data sensitivity	
Risk rating (Low / Medium / High)	
Decision (Approve / Conditions / Reject)	
Conditions or required controls	
Reviewer	
Decision date	
Re-review date	

Govern AI like an auditor would.

This form is free to use, copy, and adapt. If you want help operationalizing AI governance, from intake to control testing, that is what NexTrend Advisory does.

NexTrend.

Usman Hamid, MBA . Founder, NexTrend Advisory LLC

8+ Years in Audit & Risk . Big 4 Alum . Fortune 500 Internal Audit

nexttrendadvisory.com