

AI-Augmented Risk Use Case Canvas™

Define where AI strengthens risk detection, judgment, and response.

What decision is currently being made too late or with insufficient signal?

Early Signals (3–7 measurable indicators)

What is the cost of delay?

Are these leading or lagging?

What specific decision will this AI use case improve?

Data Sources
Structured Data

- ALM
- CI/CD
- Incident Logs
- Financials
- Resource Allocations
- Other

How often is this decision made? (Daily / Sprint / Monthly / Quarterly)

Unstructured Data

- Meeting Notes
- Status Reports
- Risk Logs
- Email
- Chat
- Vendor Updates
- Other

Risk Type

- Delivery
- Dependency
- Quality
- Compliance
- Vendor
- Adoption
- Other

Risk Appetite / Threshold

At what probability does this require action?

Signal Readiness Check

Data completeness %:

Data latency:

Known blind spots:

At what impact level does this escalate?

Who defines tolerance?

Known bias in reporting:

Business & Value Justification

Current Pain

% surprise risks: %
Average detection delay:
Late-stage re-plan frequency

Expected Impact

Faster detection by %
Reduction in rework by %
Hours saved per month
Financial impact estimate

AI Role Definition**AI Function (Select All That Apply)**

- Detect anomalies
- Classify risk types
- Summarize risk posture
- Forecast probability
- Recommend mitigation
- Monitor drift

Model Type

- Rules-based
- Predictive ML
- Generative AI
- Hybrid

Action Playbook

When triggered, what happens?

Who is notified?

What SLA for response?

Is mitigation automated or human-led?

Human-in-the-Loop

Who reviews AI output?

At what threshold?

Escalation path?

Override authority

Governance & Responsible Use**Policies**

- Consent confirmed
- Retention defined
- RBAC applied
- Audit trail enabled
- Data classification reviewed

Model Monitoring

Performance reviewed every:

Drift indicators

Recalibration owner

Kill switch condition:

Human Judgement Reflection

What might this AI miss?

Where could overconfidence create new risk?

What unintended behavior might this incentivize?

Where must human intuition override automation?

Thin Slice Implementation Plan

- Team
- ART/Program
- Portfolio
- Value Stream
- Risk Class
- Other

Pilot Duration

weeks

Success Metrics**Scale Decision Date**

****AI-Augmented Risk Use Case Canvas™**

Companion Guide**

Overview

This template helps you think clearly about how to use AI to improve risk management.

Instead of jumping straight into tools, this canvas helps you:

- Understand the decision you are trying to improve
- Identify early warning signals
- Choose the right data and AI approach
- Define how humans stay involved
- Plan a small, safe way to test your idea

The goal is simple:

Use AI to detect risks earlier and make better decisions—without losing human judgment.

This canvas works best for:

- Program and project managers
- PMOs
- Risk managers
- Team leaders working on complex initiatives

How to Fill Out Each Section

1. Decision & Risk Context

This section defines the problem you are solving.

Answer these questions:

- What decision is happening too late?
- What happens when it's late? (cost, delays, rework)
- What decision do you want to improve with AI?
- How often is this decision made?

Then select the type of risk (delivery, vendor, compliance, etc.).

Finally, define:

- When should action be taken?
- When should it escalate?
- Who decides what is “too much risk”?

 **Use this section to get clear on the problem before thinking about AI.**

2. Early Signals


List 3 to 7 signals that could warn you about risk early.

Examples:

- Missed milestones
- Increased defect rates
- Vendor delays
- Negative tone in status updates

For each signal, ask:

- Is this a leading signal (early warning)?
- Or a lagging signal (after the problem happens)?

 **Focus on signals that help you see problems sooner, not later.**

3. Data Sources

Identify where your signals will come from.

Split into two types:

Structured Data

- Tools and systems (ALM, CI/CD, financials)

Unstructured Data

- Text and conversations (emails, meeting notes, chat)

👉 **Good AI depends on good data. Be honest about what you actually have.**

4. Signal Readiness Check

Check how reliable your data is.

Fill in:

- Data completeness (how much data you actually have)
- Data latency (how fresh it is)
- Blind spots (what you can't see)
- Bias (where reporting may not be accurate)

👉 **If your data is weak, your AI results will be weak. Fix this first if needed.**

5. Business & Value Justification

This section explains why this use case matters.

Current Pain

- % of risks that surprise you
- How long it takes to detect risks
- How often you must re-plan late

Expected Impact

- Faster detection
- Less rework
- Time saved
- Financial value

👉 **Keep it simple. You are building a case for why this is worth doing.**

6. AI Role Definition

Define what the AI will actually do.

Select functions like:

- Detect patterns
- Classify risks
- Forecast problems
- Recommend actions

Then choose the type of AI:

- Rules-based (simple logic)
- Predictive (patterns from data)
- Generative (text analysis, summaries)
- Hybrid (mix of approaches)

👉 **Don't overcomplicate this. Start with what you actually need.**

7. Action Playbook

Describe what happens when a risk is detected.

Define:

- What triggers action
- Who gets notified
- How fast they must respond (SLA)
- Whether the response is automated or human-led

👉 **AI without action is just noise. This is where value happens.**

8. Human-in-the-Loop

Decide where humans stay involved.

Answer:

- Who reviews AI output?
- At what threshold?

- What is the escalation path?
- Who can override the AI?

 **This is critical. AI supports decisions—humans own them.**

9. Governance & Responsible Use

Make sure your use of AI is safe and responsible.

Check:

- Consent and data use rules
- Data retention
- Access control (who can see what)
- Audit tracking
- Data classification

Also define:

- How often the model is reviewed
- How you detect drift
- Who updates the model
- When to shut it down (kill switch)

 **This protects your organization from risk created by AI itself.**

10. Human Judgment Reflection

Pause and think critically.

Ask:

- What might AI miss?
- Where could we trust it too much?
- What bad behaviors could this create?
- When must humans step in?

👉 **This is where you close the Capability–Judgment Gap.**

11. Thin Slice Implementation Plan

Start small.

Define:

- Where you will test this (team, program, etc.)
- How long the pilot will run
- What success looks like
- When you will decide to scale

👉 **Do not try to roll this out everywhere at once. Test, learn, improve.**

Final Thought

This canvas is not about building perfect AI.

It's about:

- Making better decisions
- Seeing risk earlier
- Keeping humans in control

Start small. Learn fast. Improve over time.