SafeAI Engine™

safeaiengine.com

RESPONSIBLE AI ADOPTION STARTER KIT

Zero Essentials

> *The three questions every AI use case must answer before it is approved, deployed, or scaled.*

01	02	03
OWNERSHIP	**RISK**	**CONTROLS**
Who owns it?	*What could go wrong?*	*What stops it?*

Approve only when all three are clear.

You cannot govern, scale, or trust AI until three questions are clear: who owns it, what could go wrong, and what stops it. SafeAI Engine™ Zero Essentials is the baseline every organization must clear before AI is approved, deployed, or scaled. It is not the ceiling. It is the floor.

PURPOSE

Three questions. Every AI use case. No exceptions.

SafeAI Engine™ Zero Essentials is built on one principle: governance does not require complexity. It requires clarity. Before any AI use case is approved, deployed, or scaled, three questions must be answered.

Who owns it. What could go wrong. What stops it.

> *Zero Essentials principle: Three questions answered consistently across every AI use case are more powerful than a complex framework applied inconsistently. Start here. Build upward.*

Designed for high-trust environments

SafeAI Engine™ Zero Essentials is designed for organizations where AI may affect customers, employees, citizens, regulated data, public trust, financial outcomes, operational resilience, or executive accountability.

For public-sector environments, the baseline supports accountable AI use, transparency, data stewardship, and responsible adoption in services or decisions that may affect citizens.

THE SAFEAI ENGINE™ ZERO ESSENTIALS MODEL

01	02	03
OWNERSHIP	RISK	CONTROLS
Who owns it?	*What could go wrong?*	*What stops it?*

Approve only when all three are clear.

Why three pillars

Ownership, Risk, and Controls mirror how regulators, boards, and risk functions already think. Every mature governance framework reduces to some version of these three questions. Purpose and data matter; they are inputs to risk assessment, not separate pillars.

Three pillars also scale. A low-risk internal tool clears the same three gates as a high-risk customer-facing system. The depth of the answers differs. The structure does not.

HOW TO USE THIS STARTER KIT

One template per pillar. Complete all three.

Use this kit once for every AI use case, before you design, build, buy, pilot, deploy, or materially change an AI system. The goal is not paperwork. The goal is clarity.

Who should complete this?

1. A business owner accountable for the outcome and the risk decision.
2. A technical owner who understands how the system works.
3. Someone from risk, security, privacy, legal, or compliance for higher risk use cases.

How to use this in 30 minutes

5 min: Define ownership, use case purpose, and expected outcome.

10 min: Identify risk, data involved, people impacted, and vendor exposure.

10 min: Agree controls, monitoring indicators, evidence required, and escalation path.

5 min: Assign next review date and approval path.

Three steps

1. Complete Template 1: Ownership. Establish who is accountable, what the system does, and what outcome is expected.
2. Complete Template 2: Risk. Identify what could go wrong, how severe the impact could be, and what data is involved.
3. Complete Template 3: Controls. Agree on human oversight, monitoring, evidence, escalation, and the trigger to pause or roll back.

> *If you cannot complete a template, that is the signal. A use case without a clear owner, understood risk, or defined controls is not ready to scale.*

PART 1 · TEMPLATES

Complete one template per pillar for every AI use case. Keep answers plain-language, practical, and honest. If a team cannot answer a field, that is useful information; it means the use case needs more clarity before it scales.

The goal is proportional control: lighter review for low-risk use cases, deeper review for high-impact, regulated, public-facing, or externally facing uses.

01 Ownership

Establish accountability before anything else.

FIELD	YOUR NOTES
USE CASE NAME	
BUSINESS OWNER *Person or role accountable for outcomes and the risk acceptance decision.*	
TECHNICAL OWNER *Person or team responsible for build, integration, and operation.*	
INTENDED USERS *Employees / customers / citizens / vendors / executives / other.*	
WHAT IT DOES *One sentence in plain language.*	
OUTCOME EXPECTED *What value, efficiency, or improvement is this meant to create?*	
SUCCESS MEASURE *Baseline: current performance. Target: expected improvement. Review date: when results will be checked.*	
DECISION TYPE *Information only / recommendation / approval support / automatic decision / automatic action.*	
GO-LIVE STATUS *Idea / pilot / in use / scaling / retiring.*	

02 Risk

Understand what could go wrong before it does.

FIELD	YOUR NOTES
People impacted *Employees / customers / citizens / vendors / none / unsure.*	
Worst realistic impact *Low / Medium / High / Critical: add a short explanation.*	
Data involved *Key data sources, sensitivity, and whether personal or regulated data is used.*	
Data permission *Are we allowed to use data this way? Yes / No / Unsure.*	
External AI or vendor *Yes / No. If yes, name the provider. Does data leave the organization?*	
Vendor terms reviewed *Yes / No / Unsure. Note restrictions on data use, retention, training, confidentiality, output ownership, subcontractors, or audit rights.*	
Output explainability *Can we explain why the AI produced this output? Yes / No / Unsure.*	
Overall risk level *Low / Medium / High / Critical: one line explaining why.*	

> *Pause point: If 'Unsure' appears often, or overall risk is High or Critical, pause and request deeper review before the use case goes live or scales.*

03 Controls

Define what keeps the AI controlled, monitored, and safe to use.

FIELD	YOUR NOTES
Human oversight *Who can review, approve, or override AI outputs, and in which situations.*	
When AI must not decide *Cases where a human must decide instead of the AI.*	
Testing before go-live *Has this been tested? Yes / No / In progress.*	
How users are informed *What staff, customers, or stakeholders are told about AI involvement.*	
What we will monitor *2–3 indicators: complaints, override rate, error rate, drift, or incident reports.*	
Review frequency *Daily / Weekly / Monthly / Quarterly / After major change; include who is responsible.*	
Next review date or trigger *Date of next review, or event that requires review: model change, vendor change, data change, incident, complaint pattern, or expanded use.*	
Required evidence *Approval record, risk acceptance decision, testing records, prompts or input examples, outputs, monitoring results, issue logs, and review history.*	
Pause or rollback trigger *Clear condition that stops, pauses, or rolls back the use case.*	
Escalation path *Who is notified when the pause or rollback trigger is met.*	

PART 2 · DECISION GUIDANCE

When to pause or escalate

Pause the use case and request deeper review if any of the following apply:

- There is no named business owner or technical owner.
- The use case affects customer, employee, citizen, credit, financial, health, legal, employment, eligibility, or access-related decisions.
- Personal, sensitive, confidential, regulated, or third-party data is used, and the permission basis is unclear.
- Vendor or platform terms have not been reviewed for data use, retention, training, output ownership, or audit rights.
- The AI output could materially affect people, money, rights, safety, compliance, public trust, or business operations.
- The system takes automatic action without human review.
- There is no monitoring plan or review cadence.
- There is no pause, rollback, or escalation trigger.
- The overall risk level is High or Critical and no deeper review has been completed.

What comes next?

Zero Essentials is the first control layer. Mature AI adoption requires the same three questions to be applied across the AI portfolio, operating model, vendors, controls, reporting, and assurance.

- Maintain an AI inventory with Ownership, Risk, and Controls records for every use case.
- Define AI risk tiers to determine which use cases require deeper review, legal sign-off, or board visibility.
- Establish approval and escalation paths for medium and high-risk use cases.
- Create data, vendor, model, and security standards that sit above the Zero Essentials baseline.
- Monitor performance, incidents, complaints, and control effectiveness across the AI portfolio.
- Report material AI risks, adoption progress, and measurable outcomes to leadership and the board.

> *Core idea: Three questions answered consistently across every AI use case are more powerful than a complex framework applied inconsistently. Start here. Build upward.*