

# About me

---

- 30+ years in international IT experience
- Educator in Humber College, Toronto
- CEO of Light Ring Technology Inc, A Canadian IT Solution provider since 2012
- Cofounder of Col8.ai, an AI startup in Knowledge Management and AI Document processing.

The infographic features the company name "Light Ring Technology" at the top center. Below it are five icons, each representing a key business metric: Canadian Business (maple leaf), Engineering Team (gears), Global Presence (globe), AI Software Systems (computer monitor), and Local Projects (gears). Each icon is accompanied by a brief description of that metric.

**Light Ring Technology**

- Canadian Business**  
Operating as a Canadian business since 2012.
- Engineering Team**  
Employs over 50 software and AI engineers.
- Global Presence**  
Has operations in North America, the Middle East, and India.
- AI Software Systems**  
Focuses on delivering AI-powered software systems for web and mobile.
- Local Projects**  
Engages in local projects with the Iraqi public and private sector.

The screenshot below shows the company's website. The header includes the "LIGHT RING TECHNOLOGY" logo, navigation links for "Services", "Company", "Case Studies", and "English", and a "Book a Consultation" button. The main content area features the headline "Transform Operations with AI & Custom Software" and a sub-headline "AI, automation, and custom software for smarter, faster organizations." Below this are two buttons: "Book a Consultation" and "Explore Case Studies". The footer lists service categories: "AI Solutions", "Data & Analytics", "Process Automation", "Custom Software", and "UI/UX Design".

# Designing the Business Process for AI-Powered Solutions

From Model to Production with Confidence  
Light Ring Technology Inc. 2026

# Foundation of Business Process Design

Design business process for AI

# The Iceberg of AI

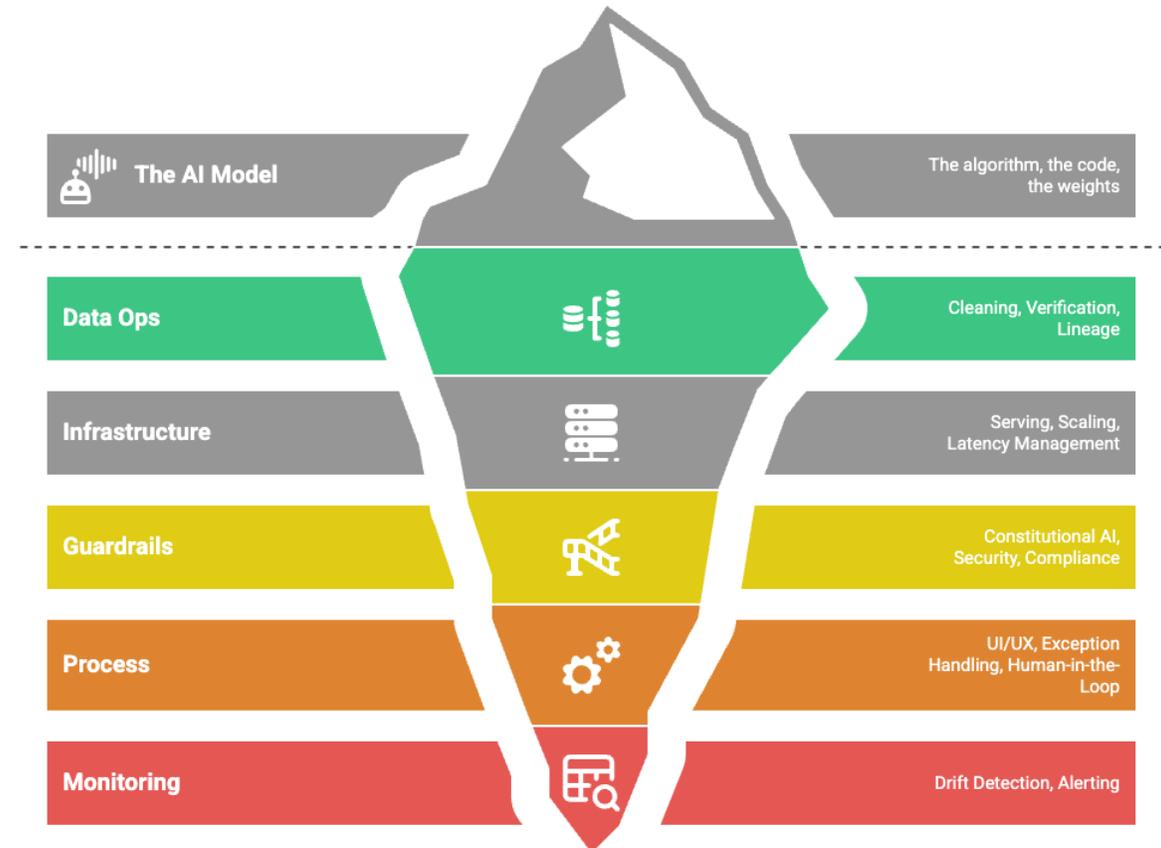
We spend our resource on the tip of the iceberg. But it is the ice below the water that sinks the ship.

- **The Model is the core math:** It is a static mathematical artifact that represent the captured learning of AI
- **The System is the Asset:** It is the living infrastructure that turns predictions into business actions.
- **The Failure Stat:** "85% of AI projects fail or get delayed not because the math is wrong, but because the submerged infrastructure is inadequate."

## The Question:

How to design effective AI powered Systems that deliver business value and meet customers expectations and trust

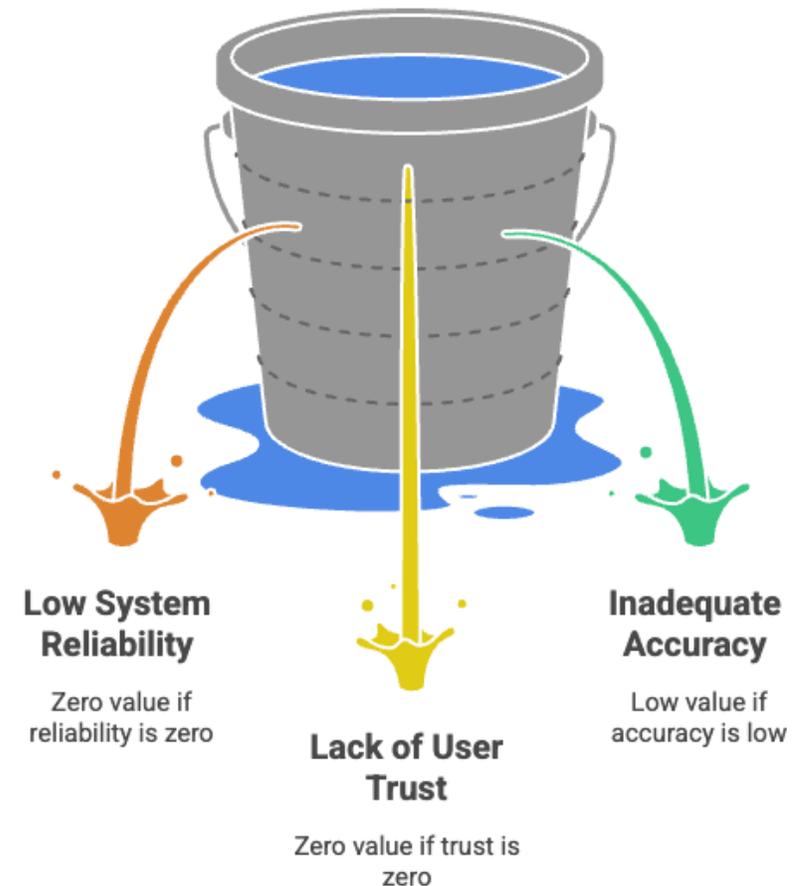
The AI Reality: The Model is Just the Tip



# The AI Business Viability Formula

$$\text{Business Value} = \text{AI Accuracy} * \text{System Reliability} * \text{User Trust}$$

- **It is Multiplication, Not Addition:** If any single factor is **zero**, the total Business Value is **zero**.
- **Scenario A (The Lab Experiment):**
  - 99% Accuracy \* 10% Reliability \* 0% Trust = **0 Value**.
- **Scenario B (The Viable Product):**
  - 85% Accuracy \* 99% Reliability \* 90% Trust = **High Business Value**.



# An AI model is not a solution; a robust business system is

## The Goal:

Not just accuracy, but operational credibility, trust, and resilience.

## Key Pillars:

- Pillar 1: Human-Centric Process Design
- Pillar 2: Designing for AI's Probabilistic Nature
- Pillar 3: The Lifecycle for Continuous Improvement
- Pillar 4: Governance, Risk, and Compliance (GRC)



### Human-Centric Process Design

Focuses on designing processes that prioritize human needs and interactions.



### Designing for AI's Probabilistic Nature

Addresses the challenges of designing systems that account for AI's inherent uncertainties.



### MLOps Lifecycle for Continuous Improvement

Emphasizes the continuous improvement of AI models through iterative processes.



### Governance, Risk, and Compliance

Ensures that AI systems adhere to ethical standards and regulatory requirements.

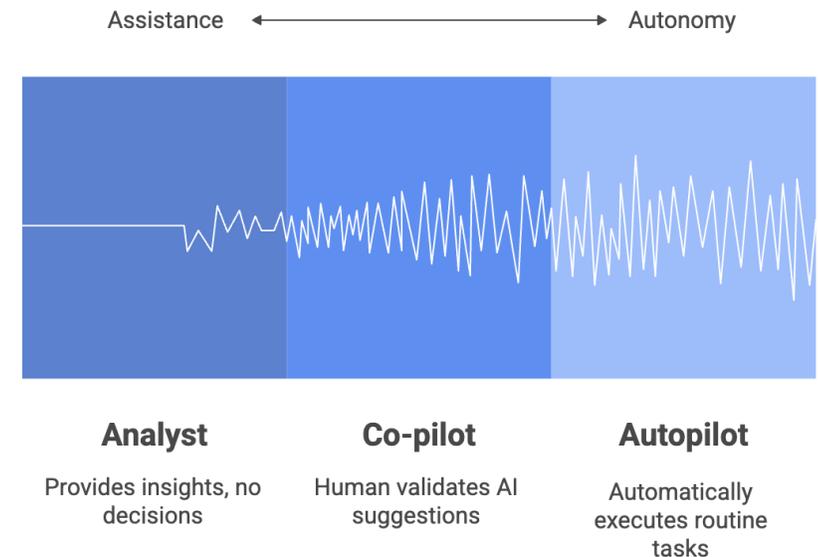
# Start with the Human: Defining the Role of AI

**Important:** Before any code, define the interaction. What is the role the AI?

## Potential Roles:

- Analyst: (Provides insights from vast data; no direct decision-making).
  - Example: AI identifies a new trend in product defects.
- Co-pilot: (Recommends, drafts, analyzes; human makes the final call).
  - Example: AI suggests a risk score, officer validates.
- Autopilot: (Handles routine tasks automatically under strict conditions).
  - Example: AI auto-approves applications with >99% confidence.

This choice is the single most important factor for safety and process design

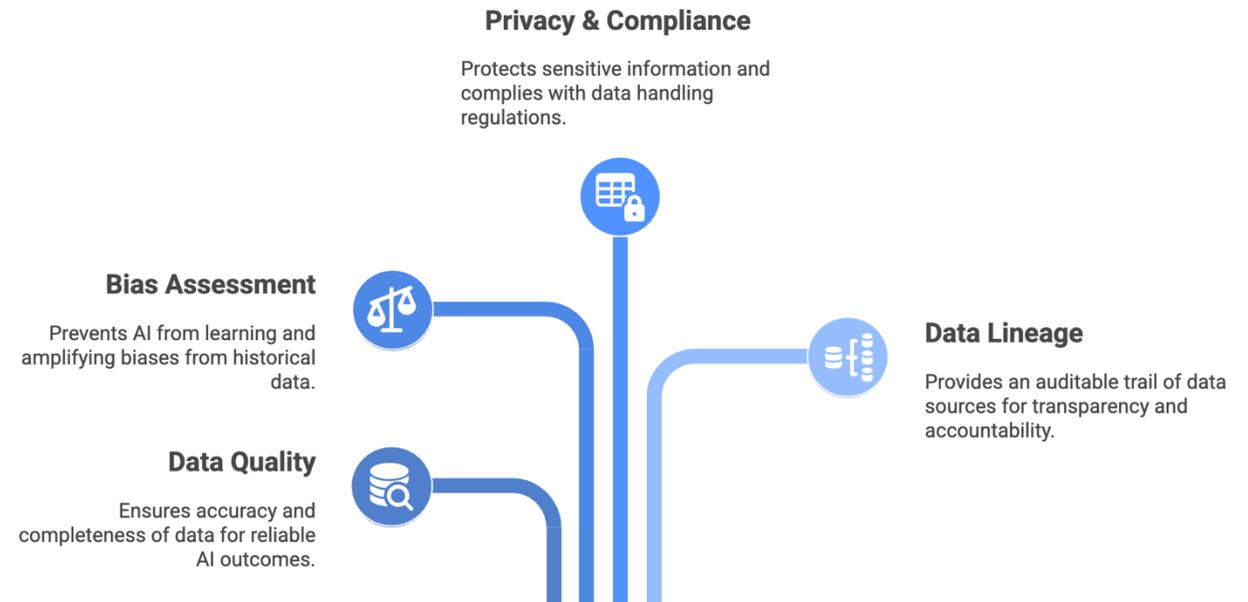


# Prerequisite: Data Governance and Readiness

An AI powered process accuracy is a reflection of its input data

## Questions:

- **Quality:** Is our data clean, accurate, and complete?
- **Bias:** Does our historical data contain biases that the AI will learn and amplify?
- **Privacy & Compliance:** Have we scrubbed sensitive information (PII)? Does our data handling comply with data privacy regulations
- **Lineage:** Do we have an auditable trail of where our data comes from?



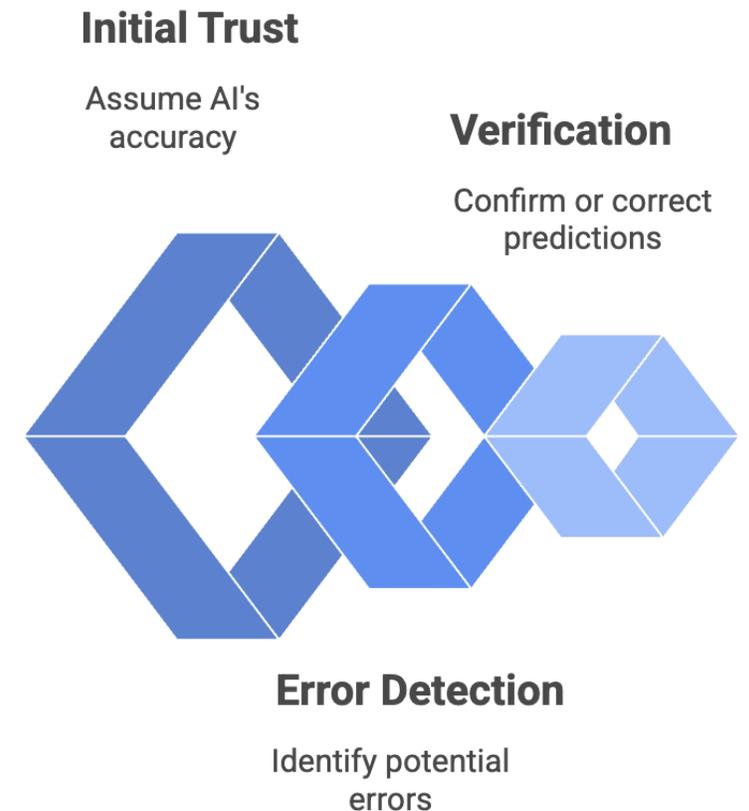
# The Guiding Principle: "Trust, but Verify"

Just like in any Business System, This principle will guide all subsequent design choices.

- AI is probabilistic; it makes highly educated guesses, not deterministic calculations.
- Our process must assume errors can and will happen and have guardrails in place to detect and handle them

## Example: Document Data Extraction

- **Trust:** We **trust** the AI to read and extract key data (e.g., dates, serial numbers, lab and test names) from 99% of submitted forms, saving thousands of manual hours.
- **Verify:** We **verify** by automatically cross-referencing critical data points against a "source of truth." For instance, an extracted "Testing Lab ID" is instantly checked against our official database of accredited labs. A mismatch, regardless of AI confidence, automatically flags the document for human review.



# Interaction Patterns: Handling Uncertainty & Latency

The Key is to let users know, let them choose, and learn from their choices.

- **The "Ambiguity Resolution" Pattern:**

- *Problem:* The AI is 70% sure.
- *Bad UI:* Show an error or guess wrong.
- *AI UI:* Ask the user. "I found two dates: Jan 1st and Jan 10th. Which is the invoice date?" (Transform uncertainty into a user choice).

- **The "Citation" Pattern (Trust Anchor):**

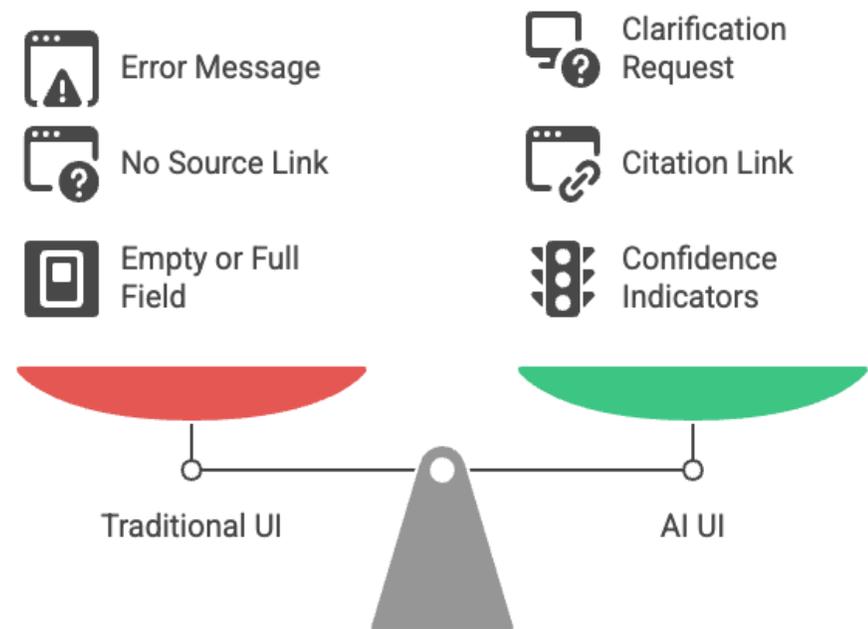
- Never display a generated summary without a link to the source. Even better explain the reasoning where needed.
- **UI Feature:** When the user hovers over an AI insight, the system should highlight the *exact* paragraph in the original document that generated that insight.

- **Confidence Signaling:**

- Use visual cues (Green/Amber/Red indicators) next to data fields to suggest where the human needs to pay attention.

- **Pre-Processing / Background Processing Pattern**

- Process content ahead of user interaction where possible
- "Fire and Forget": Let users start a complex analysis and leave the page. Notify them (Slack/Email) when the 'Agent' is done



# Implementing Explainable AI (XAI)

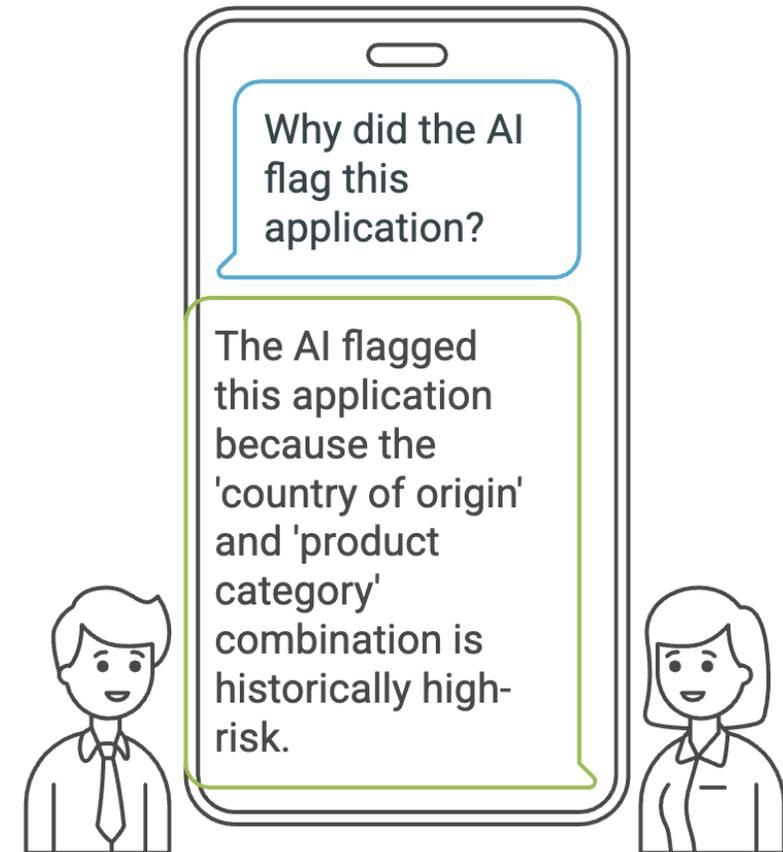
"The AI said so" is not an acceptable answer for an auditor, the system must be able to answer why

Every AI recommendation comes with a clear, human-understandable explanation. This is essential for building trust with our team and creating a rock-solid audit trail.

## Ex. For Document Analysis

- **Instead of:** A vague "non-compliant" status, the AI **highlights the specific clause or data point** in the document that violates our regulations, so the officer knows exactly where the problem is.
- This transparency empowers officers to make **faster, more confident decisions**, knowing exactly what the AI is seeing and why.

Designing for Explainable AI ensures the business is always in control and can defend every single decision.



# Designing for Operational Realities

Tuning Business Processes to the Realities of AI

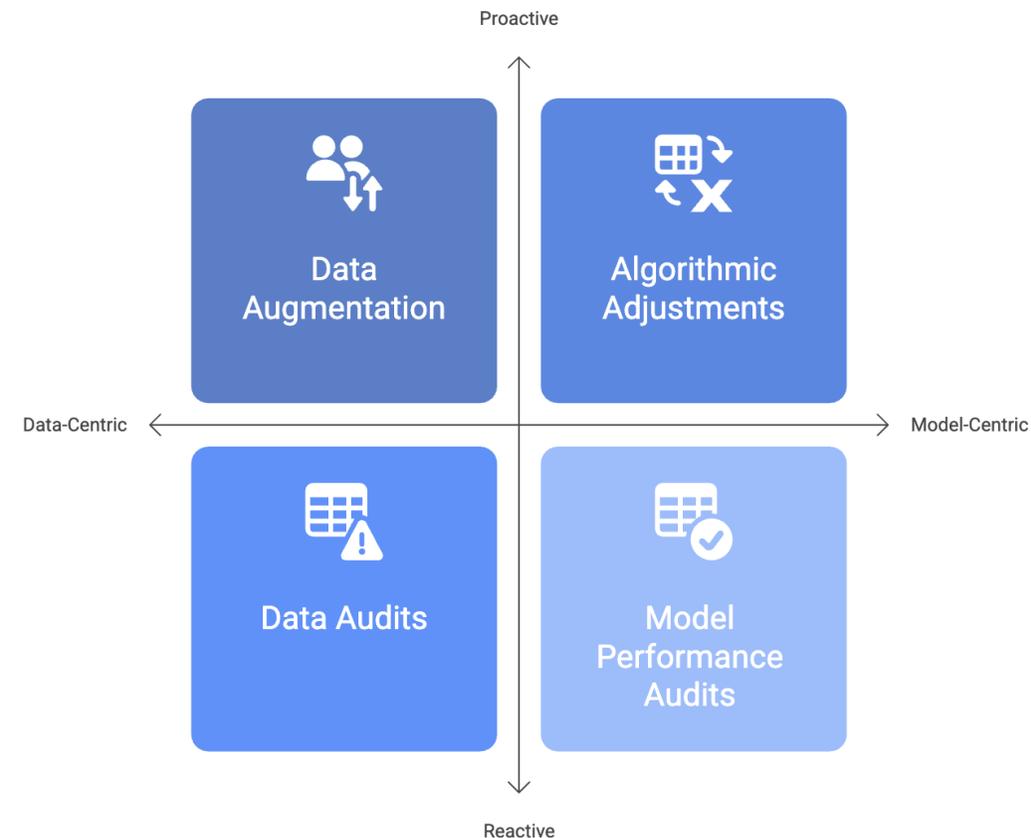
# Handling Inaccuracies Part 1: Handling Bias in AI

Bias happens when AI learns incorrect assumptions because of data inaccuracies. It is not a technical problem, rather is it a data and human problem that AI learns and amplify

## Example: Bias in Product Compliance Data

- If we historically inspected products from USA ten times more than from UK, our data is full of non-compliance examples from USA. AI trained on this data, would incorrectly learn that US products are more non-compliant than UK products.
- **The Result:** The AI will learn to associate USA with higher risk, leading to unfair scrutiny, even if its products are now fully compliant. It's not because the AI is "biased" against USA, but because our historical data gave it an unbalanced education.

N.B: Bias also happens with humans. We often call it a Stereotype



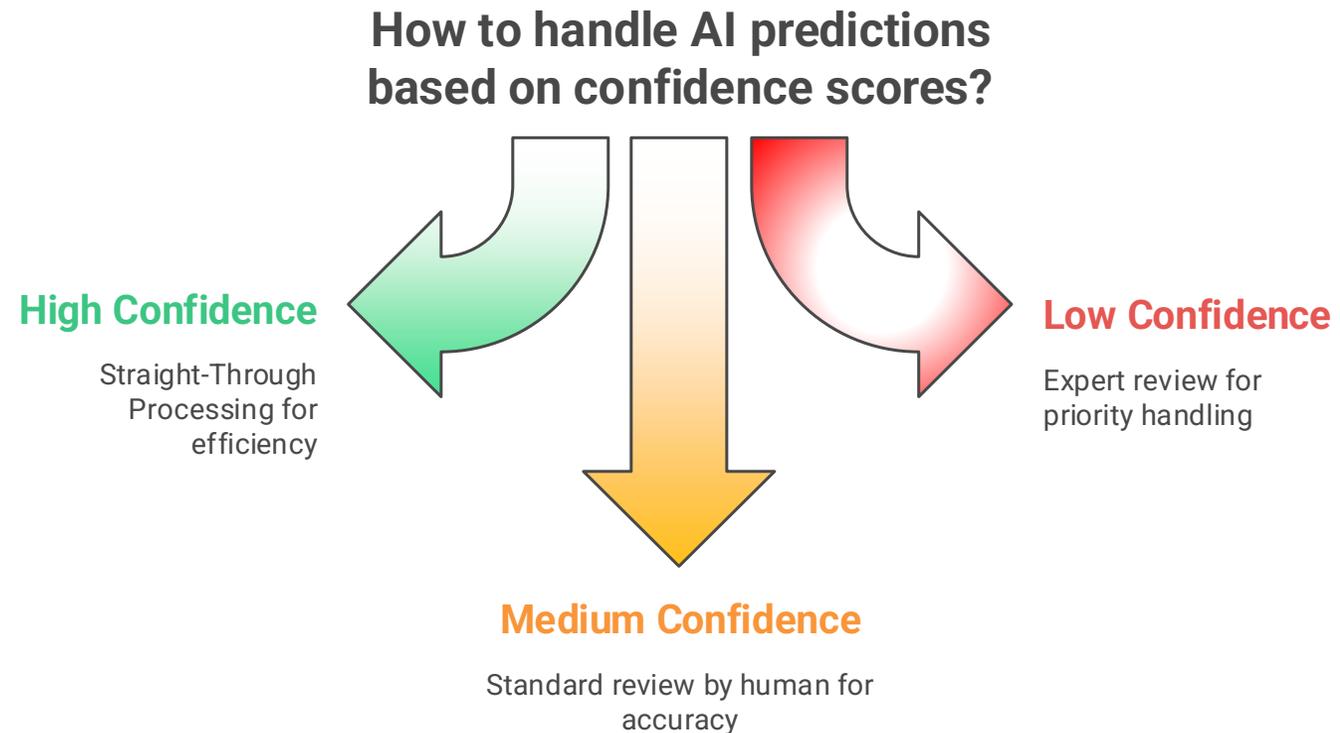
# Handling Inaccuracies Part 2: The "Confidence Score"

Instead of a single process pipeline, design the process with multiple branches

## Example: Creating a three-channel business process

The business process branches based on this score:

- **High Confidence (>95%):** Straight-Through Processing (Autopilot).
- **Medium Confidence (70-95%):** Route to human for standard review (Co-pilot).
- **Low Confidence (<70%):** Escalate to a human expert for priority review.



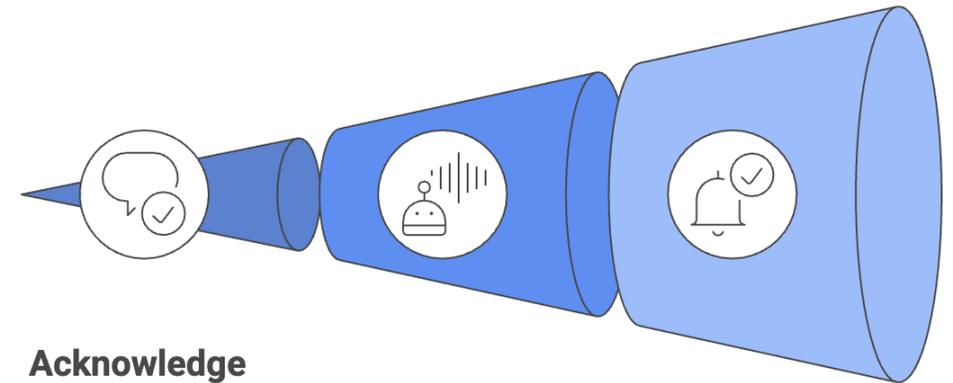
# Handling AI Latency: Designing for a "Patient" Process

Complex AI queries can take seconds or minutes, not milliseconds. Standard user interfaces will fail.

Instead of keeping the user waiting a spinning wheel, design for asynchronous workflow Process

- **The Right Way: Asynchronous Workflows**
  - **Acknowledge & Notify:** The system immediately confirms receipt of the request ("Your analysis is underway..").
  - **Process in Background:** The AI performs its work without blocking the user.
  - **Callback/Notification:** The system notifies the user via email, dashboard update, or alert when the result is ready.
  - **Progress Monitoring:** Allow the user to track progress
- This turns a "slow system" into an "efficient background process."

## Transforming AI Queries into Efficient Processes



### Acknowledge Request

System confirms receipt of the query

### Process in Background

AI performs analysis without blocking user

### Notify User

User is informed when results are ready



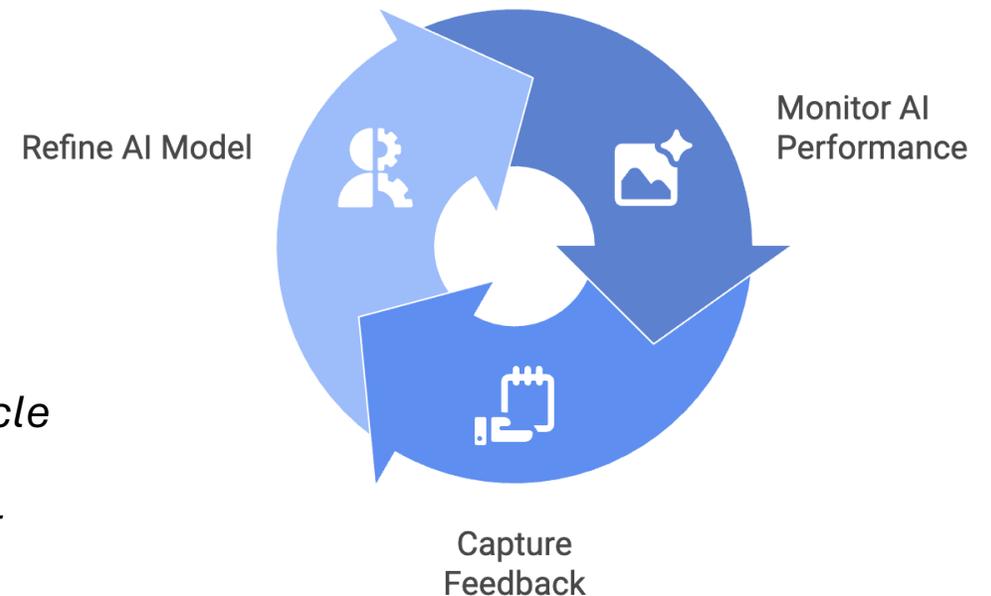
# Continuous Improvement Lifecycle of AI Models

Unlike static software systems, an AI model is a dynamic capability that must evolve with the business and data evolution

## Example: A New Safety Regulation is Introduced

- **Monitor:** Our system detects that a new federal safety standard for e-bike batteries has been published. The current AI's knowledge is now outdated.
- **Feedback:** Officers begin manually reviewing the first few e-bike applications against the new standard. Their decisions are logged as "lessons learned."
- **Refine:** We use the text of the new regulation and the officers' initial decisions to retrain the AI. The model is now an expert on the new standard, ready to assist with future applications.

*This professional, structured approach to managing the AI lifecycle is formally known in the industry as **MLOps** (Machine Learning Operations). It's the gold standard for ensuring AI delivers lasting value.*



# Continuous Monitoring: AI “Check Engine” Light

Both technical and business performance must be monitored regularly

## What to Track:

- **Data Drift:** Is the new, real-world data significantly different from the training data?
- **Concept Drift:** Have the underlying patterns changed? (e.g., a new regulation makes old predictions obsolete).
- **Model Performance:** Is accuracy degrading? Is latency increasing?
- **Technology Drift:** Are there new algorithms or techniques that can improve system performance
- **Business KPIs:** Is the AI actually improving efficiency or quality as intended?

Automated alerts are key to catching issues before they impact the business



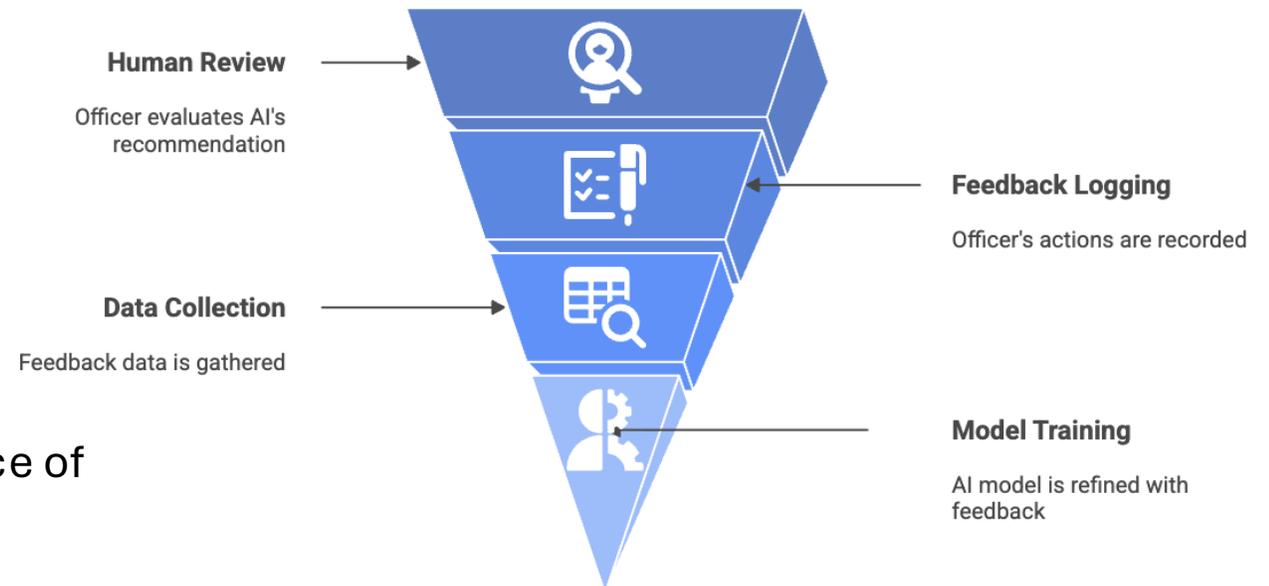
# The Human Feedback Loop: The Source of Truth

The data of collected human actions in the process is used to train a more refined AI model

## The Process Loop:

- AI makes a recommendation.
- Human officer reviews it.
- If the officer overrides or corrects the AI, this action is **logged** in a structured way. ("Officer Jane Doe rejected AI's 'compliant' status, citing 'incomplete documentation'").
- This feedback is the most valuable data source for improving the model.

Every human interaction with the AI is a valuable piece of data for future improvement.

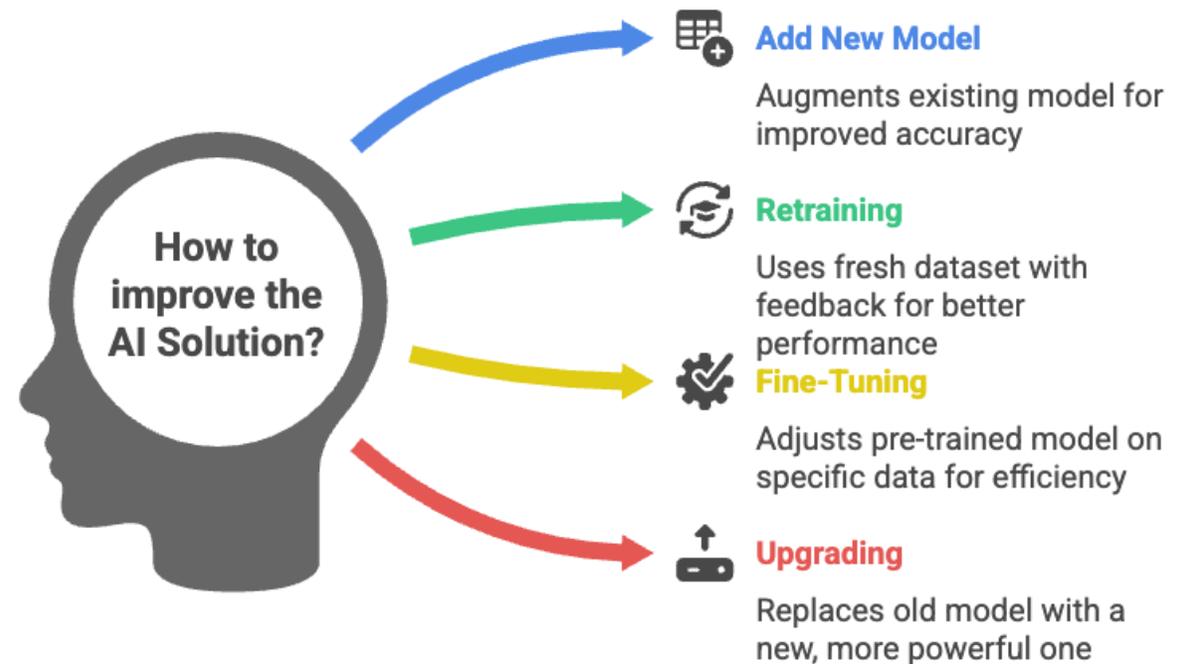


# The Human Feedback Loop: The Source of Truth

Monitoring and feedback tell the business *when* to act. Improving AI could include several choices

- **Upgrading:** Swapping out an old model entirely for a new, more powerful one (e.g., moving from GPT-4 to GPT-5).
- **Fine-Tuning:** Taking a pre-trained model and making small adjustments on a smaller, specific dataset (faster and cheaper).
- **Retraining:** Re-running the training process on a fresh dataset that includes new feedback data.
- **Adding New Model:** Add an additional model to augment the existing one and improve accuracy

This process must be planned and version-controlled, just like any software update.



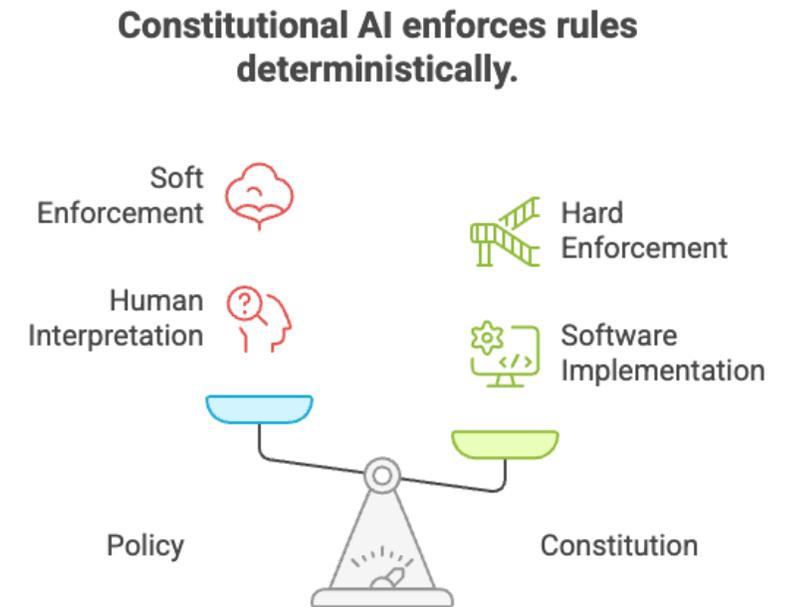
# Responsible AI: Governance, Risk and Compliance

# Constitutional AI from Policy to Enforcement

We don't trust the model to 'behave' based on a nice prompt. We enforce behavior with code. The Constitution is the sheriff that polices the AI's output.

## Key Concept

- **The Problem with Prompts:** You cannot "prompt engineer" safety into a model. If the model is probabilistic, there is always a non-zero chance it will break the rules.
- **The Definition:** Constitutional AI places a deterministic software layer *outside* the model. It is a set of "Non-Negotiable Rules" that intercept and override the model's output if necessary.
- **Policy vs. Constitution:**
  - *Policy (Governance):* "We should not give medical advice."
  - *Constitution (Code):* IF topic == "medical\_diagnosis" THEN return "I cannot answer this."
- **Business Value:** This allows regulated enterprises (Banking, Health, Gov) to deploy GenAI without fear of reputational damage or liability.



# Data is The Foundation of Trustworthy AI

An AI is like a brilliant chef—it can create amazing things, but only if it has high-quality, trusted ingredients.

Ensure the integrity of AI by focusing on three pillars of data management:

1. **Data Quality** (Are the ingredients good?)
2. **Data Governance** (Are there rules in the kitchen?)
3. **Data Provenance** (Do we know where the ingredients came from?)

An AI trained on flawed data will produce flawed results, exposing business to operational and reputational risk.

## The Foundation of Trustworthy AI



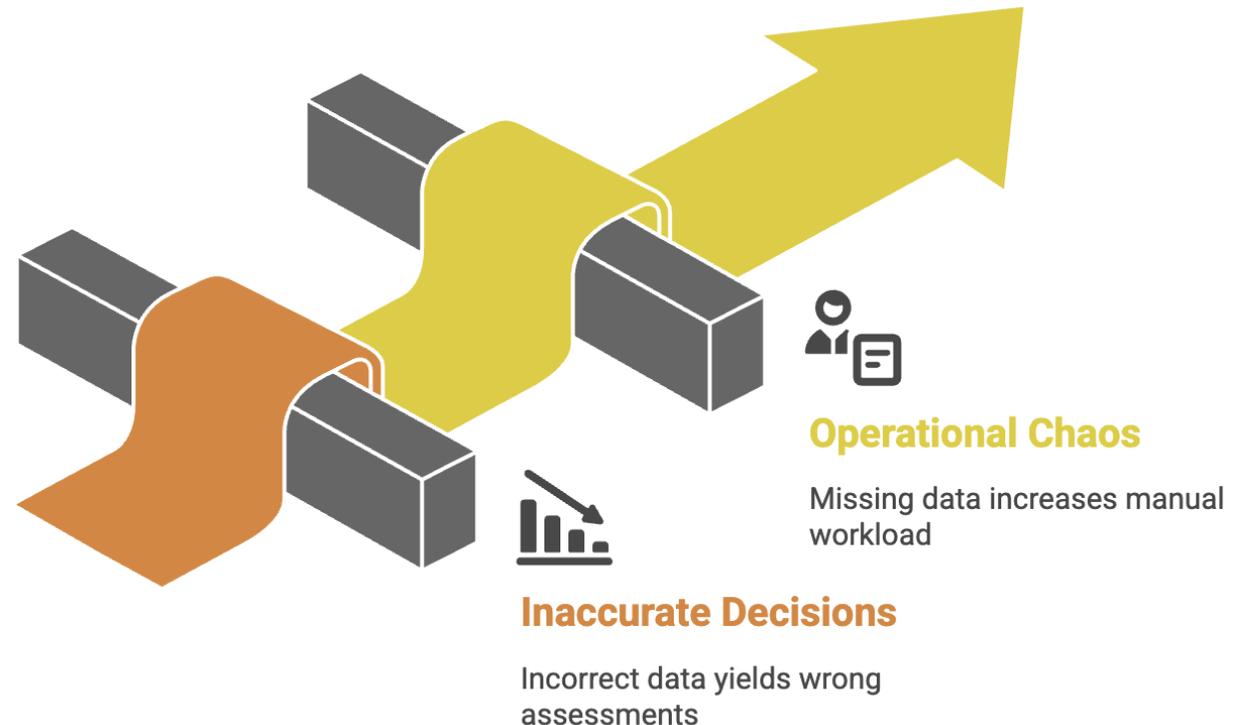
# Data Quality: Is Our Data Fit for Purpose

Is the data we feed the AI **accurate, complete, and consistent**? Can we rely on it to reflect reality?

## The Business Risk of Poor Data Quality:

- **Inaccurate Decisions:** If our database has incorrect product codes, the AI will make wrong assessments, potentially approving unsafe products or rejecting compliant ones.
- **Operational Chaos:** If 30% of application forms are missing the "country of origin," the AI's risk models become unreliable, forcing our officers to do more manual work and slowing down the entire process.

An AI trained on flawed data will produce flawed results, exposing business to operational and reputational risk.



# Data Governance: The Rules of the Road

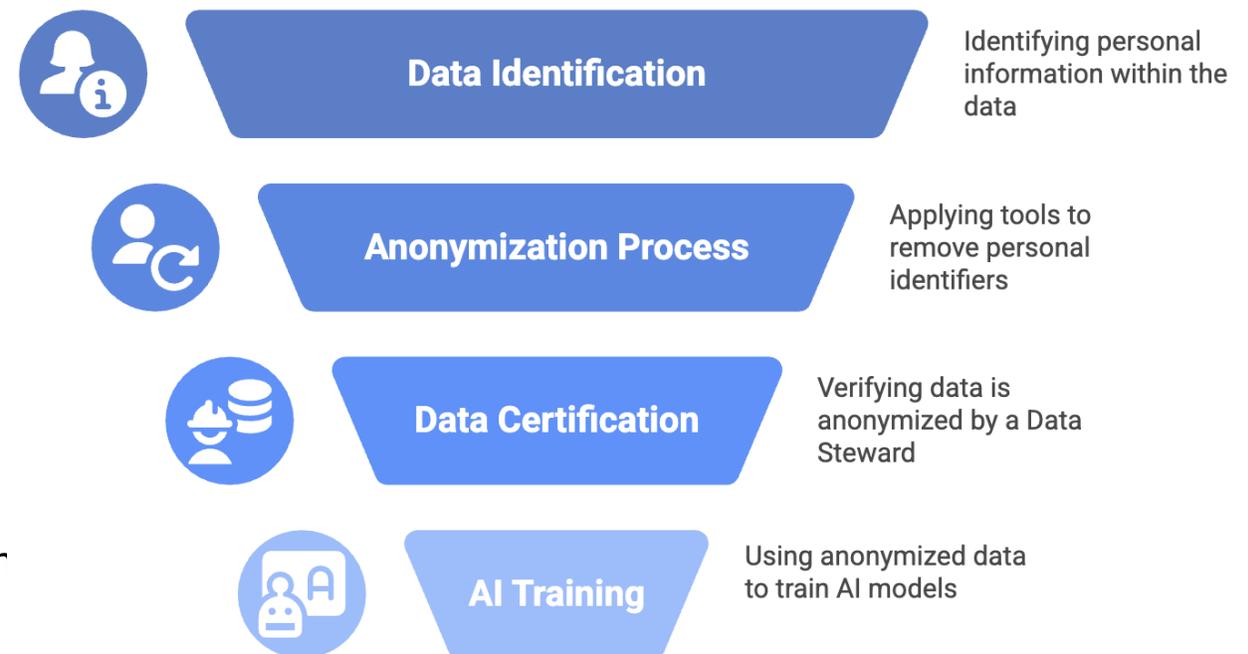
Who has **ownership and responsibility** for data? What are the clear, enforceable rules for how it can be accessed, used, and secured?

## The Business Risk of No Governance:

- **Compliance & Legal Penalties:** Without strict rules, sensitive information (like a person's contact details) could accidentally be used to train the AI, leading to a major privacy breach and violating laws or regulations.
- **Inconsistent Results:** If different teams use different data sources or rules, they will get different answers from the AI. This erodes trust and prevents the AI from being a "single source of truth."

Establishing a clear Data Governance framework ensures AI process is **secure, compliant, and consistent** across the organization, and protecting from legal and financial risk.

## Data Governance Process Funnel



# Data Provenance: The Unbreakable Audit Trail

Is it possible to trace every piece of data the AI learned from back to its **original, verifiable source**?

The scenario illustrates the critical business value of **data provenance** (an auditable data trail) when an AI-assisted decision is challenged.

## The Failure (Without Provenance)

An AI-approved toy is recalled, and during an audit, the organization cannot prove *why* the AI made its decision. The data sources are a mix of unverified and untraceable information.

- **Business Impact:** The approval process is discredited, exposing the organization to **major legal liability, financial penalties, and a catastrophic loss of public trust.**

## The Success (With Provenance)

When audited for the same recalled toy, the organization uses a "provenance log" to show the exact, verifiable, and trusted sources (e.g., official lab reports, national recall database queries) the AI used to make its approval.

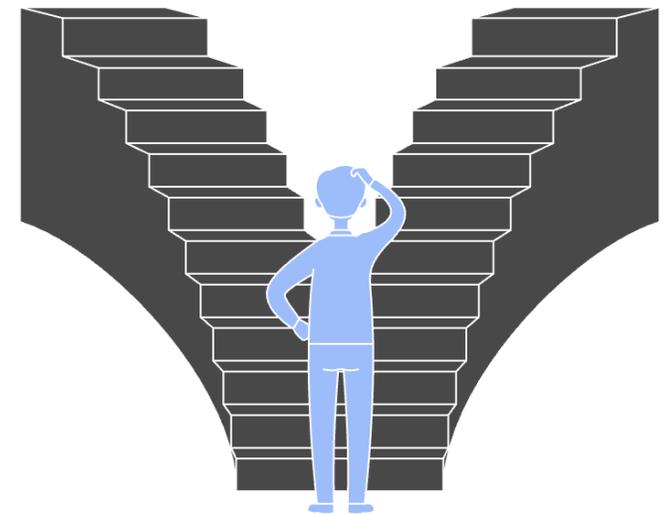
- **Business Impact:** The organization's *process* is proven to be **diligent, defensible, and robust.** This protects them from claims of negligence and upholds their institutional integrity, even if the product itself failed for other reasons.

## Establish Data Provenance

Ensures decisions are based on verifiable data, enhancing trust and defensibility.

## Use Unverifiable Sources

May lead to indefensible decisions and loss of trust due to unverifiable sources.



# Design for Auditability: Keeping a log of AI decisions

Assume every decision will be reviewed one day.

**For every AI-assisted decision, we need to log:**

- A unique transaction ID.
- The exact input data provided to the AI.
- The model's version number (e.g., risk-model:v1.2).
- The full AI response (prediction, confidence score, explanation).
- The final human action (approved, rejected, edited).
- Timestamp and user ID.

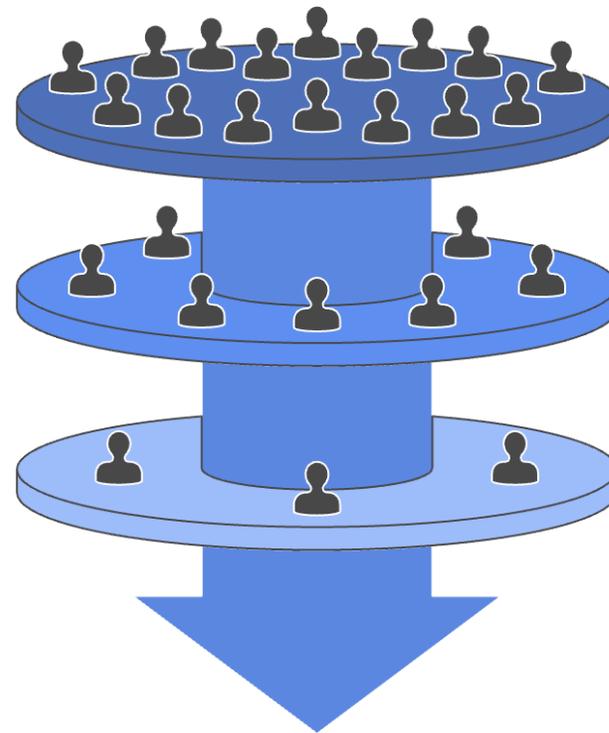


# Keep People in Mind: Change Management

An AI process is a human process. It is vital to keep people informed, engaged and empowered to manage System operations

## Key Success Factors:

- **Clear Communication:** Explain what the AI does, what it *doesn't* do, and how it will help employees.
- **Training:** Train people on how to use the new tool, interpret its outputs, and provide effective feedback.
- **Role Evolution:** Proactively redefine job roles to focus on higher-value tasks that the AI enables.



 Clear Communication

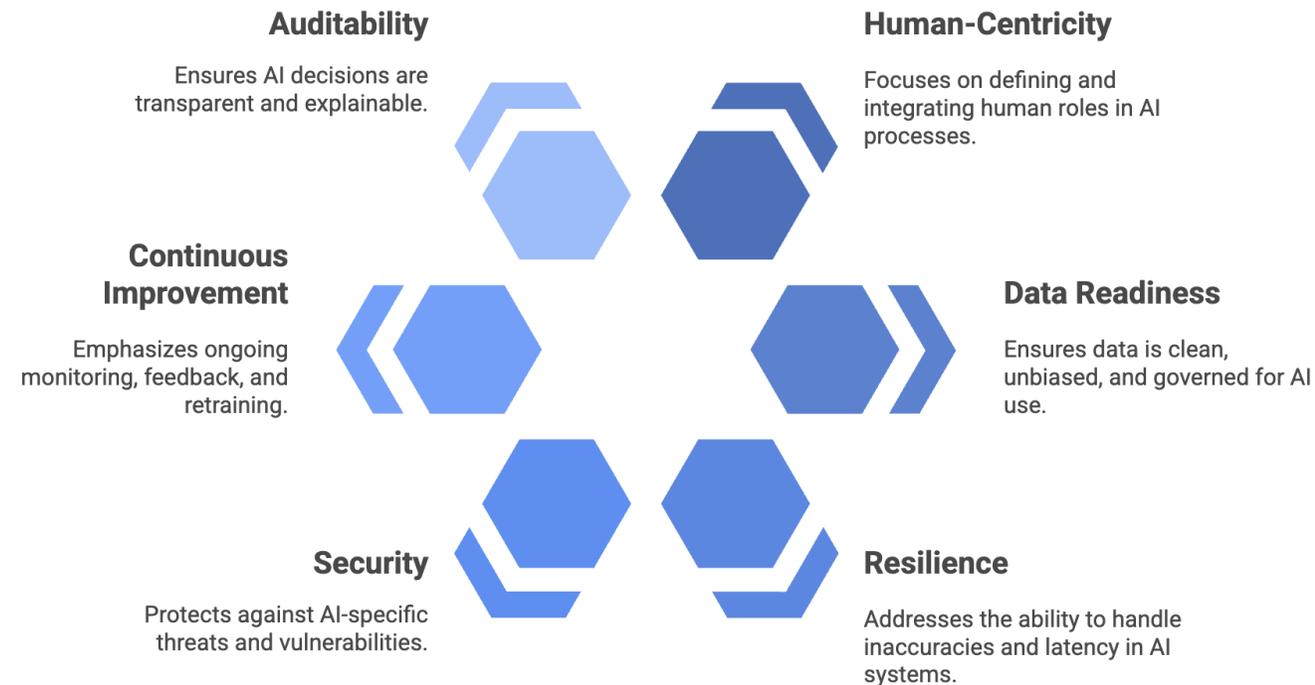
 Training

 Role Evolution

# A Checklist of Viable AI Systems

## A Checklist for Robust AI Powered Processes

- **Human-Centric:** Is the human's role clearly defined?
- **Data-Ready:** Is our data governed, clean, and unbiased?
- **Resilient:** Does the process handle inaccuracies and latency gracefully?
- **Secure & Controlled:** Have we protected against AI-specific threats? Do we have control over AI output.
- **Continuously Improving:** Is our MLOps loop (Monitor, Feedback, Retrain) in place?
- **Auditable:** Can we explain every AI decision to a third party?



# 5 Steps to Build Viable AI Systems

## 1. Define the Non-Negotiables (The Constitution)

- *Action:* Before building, identify the regulatory and business "Hard Lines" (e.g., "No medical advice," "No PII storage").
- *Output:* A deterministic set of rules that the System will enforce over the Model

## 2. Verify Data Lineage (The Foundation)

- *Action:* Audit data sources. Trace every inputs back to trusted origins
- *Output:* A "Certified" dataset that eliminates the quality risks.

## 3. Architect the "Human-in-the-Loop" (The Workflow)

- *Action:* Design the confidence thresholds. When does the AI act alone (Autopilot), and when does it ask for help (Co-pilot)?
- *Output:* A branching process diagram (Green Lane vs. Red Lane).

## 4. Operationalize the Feedback Loop (The Learning)

- *Action:* If a human overrides the AI, that data must be logged, not lost.
- *Output:* A "Golden Dataset" of edge cases used for re-training.

## 5. Monitor for Drift (The Maintenance)

- *Action:* Set up automated alerts for "Concept Drift" (Is the market changing?) and "Data Drift" (Is the input changing?).
- *Output:* A "Check Engine Light" dashboard for your AI operations.

*Establish the foundational constraints and architecture before any code is written.*

