AI Inventory and Risk Register

Template: User Guide

Tab 1: Start Here / Instructions

This tab orients anyone opening the file for the first time. Governance records are usually maintained by several people, so clear instructions prevent inconsistent entries.

What it contains

- Purpose: one or two sentences on why the workbook exists (visibility, accountability, and control over AI use).
- How to use it: a numbered sequence, such as (1) add each AI system to the Inventory, (2) log risks in the Risk Register, (3) score them using the Rating Guide, (4) review the Dashboard monthly.
- Color key: so users know which cells to edit.
    - White cells: free-text input
    - Blue cells: dropdown selection
    - Grey cells: auto-calculated, do not edit
- Version and date: for example, *Version 1.0, last updated 01 Oct 2026*, plus an owner name for the workbook itself.

Example instruction text: "Add one row per AI system. Include third-party tools such as chatbots embedded in SaaS platforms, not just models you built."

Tab 2: AI Inventory

The Inventory is the foundation of the workbook. It holds one row per AI system and should include internal models, vendor products, and AI features inside tools your teams already use.

Identification and ownership

- AI System Name: the name people actually use, such as "Customer Support Assistant" or "Resume Screening Tool."
- Business Owner: the person or department accountable for the system's outcomes, such as the Head of Customer Operations.
- Technical Owner: the team responsible for implementation and maintenance, such as the Platform Engineering team.

Every system needs both owners. If nobody is willing to be listed as the owner, that is a governance finding in itself.

Purpose and type

- Business Purpose: why the system is used, in plain language. *Example: "Drafts responses to customer support tickets for agent review."*
- AI Type (dropdown): Generative AI, Machine Learning, Agentic AI, Predictive Analytics, or Other.
- Model/Provider: internal model or third-party provider. *Example: "Third-party LLM accessed through vendor API."*

Data and users

- Data Used (dropdown, multi-value allowed): Personal, Confidential, Public, Financial, or Operational. If a system touches personal data, this flags it for privacy review.
- Users: who interacts with it: employees, customers, partners, or others.

Impact and classification

- Business Impact: what happens if the system fails or produces bad output. *Example: "Incorrect guidance could lead to customer complaints and refunds."*
- Risk Classification (dropdown): Low, Medium, or High, using the criteria in Tab 4.
- Regulatory Scope: applicable laws, regulations, and standards. *Example: "Data protection laws; internal AI acceptable use policy."*

Controls and oversight

- Security Controls: access controls, logging, encryption, and vendor security reviews.
- Human Oversight: how and when people review outputs. *Example: "Agent reviews every draft before sending."*
- Monitoring: what is tracked, such as accuracy, security events, fairness, and drift.

Lifecycle

- Review Date: when the next governance review is due. High-risk systems might be reviewed quarterly, low-risk ones annually.
- Status (dropdown): Proposed, Active, Suspended, or Retired.

Tip: Keep retired systems in the list rather than deleting them. They provide an audit trail and show that decommissioning was handled properly.

Tab 3: AI Risk Register

The Risk Register holds one row per risk, each linked to a system from the Inventory. A single system can have several risks.

Field-by-field explanation

- Risk ID: a unique code such as AI-001, AI-002. Never reuse IDs, even after a risk is closed.
- AI System (dropdown): pulled from Tab 2 so every risk traces back to a specific system.
- Risk Description: describe the risk as a cause and consequence. *Example: "The assistant may generate inaccurate refund guidance, leading to customer complaints and financial loss."*
- Risk Category (dropdown):
    - Data and Privacy
    - Security
    - Accuracy and Reliability
    - Bias and Fairness
    - Third-Party and Vendor
    - Regulatory and Compliance
- Likelihood and Impact (dropdowns): rated Low, Medium, or High using Tab 4.
- Risk Level (auto-calculated): derived from likelihood and impact so scores stay consistent.
- Existing Controls: what is already reducing the risk. *Example: "Human review of all responses; approved knowledge base."*
- Mitigation Action: what will be done next. *Example: "Add evaluation tests and escalation rules for refund topics."*
- Risk Owner: one named person responsible for the risk, such as the Customer Operations Manager.
- Target Date: the deadline for completing the mitigation.

- Residual Risk: the risk level expected *after* controls and mitigation are in place.
- Status (dropdown): Open, Monitoring, or Closed.

Worked example of one risk

Field	Entry
Risk ID	AI-001
AI System	Customer Support Assistant
Risk Description	May generate inaccurate customer guidance
Category	Accuracy / Operational
Likelihood / Impact	Medium / High
Risk Level	High
Existing Controls	Human review and response validation
Mitigation Action	Improve evaluation and escalation rules

Residual Risk	Medium
Status	Open

The goal is not to eliminate every risk. It is to understand each one, apply proportionate controls, document decisions, and keep monitoring as conditions change.

Tab 4: Risk Rating Guide

This tab exists so that two different people scoring the same risk arrive at the same answer.

Likelihood definitions

- Low: unlikely to occur in the next 12 months.
- Medium: could occur once or twice in the next 12 months.
- High: expected to occur, or has already occurred.

Impact definitions

- Low: minor inconvenience, easily corrected, no external effect.
- Medium: noticeable disruption or limited customer effect.
- High: regulatory exposure, significant financial loss, harm to individuals, or serious reputational damage.

Likelihood x Impact matrix

	Impact: Low	Impact: Medium	Impact: High
Likelihood: High	Medium	High	High
Likelihood: Medium	Low	Medium	High
Likelihood: Low	Low	Low	Medium

Inventory risk classification criteria

Define clear triggers for each class. For example, classify a system as High if it does any of the following:

- Influences decisions about hiring, lending, insurance, or healthcare
- Processes personal or sensitive data
- Interacts directly with customers without human review

Consistent criteria are one of the core best practices in AI governance. Adjust these definitions to match your organization's risk appetite.

Tab 5: Sample Entries

Seeing completed rows is often the fastest way to learn a template. This tab shows how the Inventory and Risk Register connect.

Sample inventory row

- System: Customer Support Assistant
- Type: Generative AI, third-party provider
- Data used: Personal, Confidential
- Users: Support agents
- Oversight: Agents review every draft
- Classification: High
- Status: Active

Sample linked risks

- AI-001 (Accuracy): inaccurate guidance given to customers.
- AI-002 (Data and Privacy): customer details entered into the tool may be retained by the vendor.
- AI-003 (Security): prompt injection through customer-submitted text could manipulate outputs.
- AI-004 (Third-Party): vendor outage disrupts support operations.

Notice that one system produced four risks across four categories. This is normal and shows why the register should be linked to the inventory rather than kept separately.

Label these rows clearly as samples and advise users to delete them before entering real data.

Tab 6: Governance Checklist

This tab turns the 14-point checklist into a working tool. Each item has a Yes / No / In Progress dropdown and a notes column.

Checklist items

1. Is every known AI system recorded?
2. Does every system have an owner?
3. Is the system's purpose documented?
4. Are data sources identified?
5. Has an AI risk assessment been completed?
6. Are privacy and security risks documented?
7. Are third-party providers assessed?
8. Are human oversight requirements defined?
9. Are AI outputs monitored?
10. Are incidents documented?
11. Are remediation actions tracked?
12. Are governance reviews scheduled?
13. Are employees receiving appropriate AI governance training?
14. Is evidence available for AI governance auditing?

Useful additions

- A completion percentage at the top, counting the "Yes" answers out of 14.
- An evidence link column pointing to the supporting document.

- An example note: for item 7, "Vendor security questionnaire completed 15 Sep; contract review pending."

Use the checklist during onboarding of new AI tools, during procurement, and at each periodic review.

Tab 7: Dashboard (Optional)

The Dashboard summarizes the other tabs automatically so leaders can see the state of AI governance at a glance.

Suggested panels

- Total AI systems by status: for example, 18 Active, 3 Proposed, 2 Suspended, 4 Retired.
- Systems by risk classification: a bar chart of Low, Medium, and High.
- Risks by category: shows whether exposure is concentrated in, say, data privacy or vendor risk.
- Risks by level: a count of High, Medium, and Low risks that are still open.
- Overdue actions: risks whose target date has passed and whose status is not Closed.

What to look for

- A large number of systems without owners suggests weak accountability.
- Many high risks with no target dates suggest mitigation is not being tracked.
- Reviews past their due date suggest governance has stalled.

Design and Maintenance Tips

- Dropdowns everywhere: consistent values keep filters and dashboards accurate.
- Conditional formatting: red, amber, and green on Risk Level so high risks stand out.
- Frozen header rows and filters on Tabs 2 and 3 for easier navigation.
- Review rhythm: update the Inventory whenever a new tool is adopted, and review the Risk Register at least quarterly.
- Change log: record who changed what and when, which supports audits.

Where to Go from Here

A spreadsheet is a strong starting point. As AI use grows, the same fields can feed into GRC platforms, model management systems, and compliance workflows. Start with visibility and accountability rather than waiting for a perfect platform.

GSDC
Global Skill Development Council

CERTIFIED AI GOVERNANCE PROFESSIONAL

THE CERTIFIED AI GOVERNANCE PROFESSIONAL CERTIFICATION IS BUILT AROUND AI GOVERNANCE FRAMEWORKS, RISK MANAGEMENT, REGULATORY COMPLIANCE, RESPONSIBLE AI, AND GLOBALLY RECOGNIZED INDUSTRY BEST PRACTICES.

ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

LEARNING OBJECTIVE

- Understand the fundamentals of AI governance and responsible AI adoption.
- Learn major AI governance frameworks, including ISO/IEC 42001, NIST AI RMF, and other global standards.

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org