STARTER KIT

AI Governance Committee

Charter, RACI Matrix, Meeting Templates, and a First 90 Days Checklist

Item	Detail
Version	1.0
Date	[Insert date]
Prepared by	[Your blog / organization name]
Intended audience	AI, risk, compliance, security, data, and business leaders

How to Use This Kit

Artificial intelligence is moving from experimentation into business-critical work. As it does, governance can no longer sit with a single technology or compliance team. This kit gives you the core documents needed to set up a cross-functional AI Governance Committee and get it running in its first three months.

What is inside

Part	Document	Use it to
1	Committee Charter Template	Define purpose, scope, membership, and decision rights
2	RACI Matrix	Clarify who is Responsible, Accountable, Consulted, and Informed
3	Meeting Agenda and Minutes Templates	Run consistent meetings and record decisions
4	First 90 Days Checklist	Move from kickoff to an operating governance process

Suggested order of use

Work through the parts in sequence, because each one builds on the last:

- **Start with the Charter (Part 1).** It settles authority and membership. Without it, the RACI and agendas have nothing to anchor to.
- **Adapt the RACI (Part 2)** once you know who sits on the committee.
- **Adopt the meeting templates (Part 3)** before your first meeting so decisions are documented from day one.
- **Follow the 90-day checklist (Part 4)** to track progress and keep momentum.

How to customize

Text in [square brackets] is a placeholder to replace. Example wording is provided throughout, but every organization is different. A financial institution will need deeper legal, risk, and compliance involvement than a ten-person software company, and a global enterprise may need regional or business-unit representatives. Treat this kit as a starting point, not a finished policy.

> **Tip:** Keep the committee working alongside existing enterprise risk, cybersecurity, privacy, data governance, procurement, and compliance functions. The goal is to connect them for AI decisions, not to build a separate parallel bureaucracy.

Part 1: AI Governance Committee Charter Template

The charter is the committee's founding document. It states why the committee exists, what it may decide, who sits on it, and how it reports. A clear charter prevents the two most common early failures: a committee with no real authority, and a committee that tries to review everything and becomes a bottleneck.

Each section below explains what to include, gives example wording, and points out common pitfalls.

1. Purpose and Mandate

State in two or three sentences why the committee exists. The purpose should go beyond approving or rejecting AI projects. It should describe the committee as the mechanism for responsible, accountable AI adoption.

> **Example:** "The AI Governance Committee is a cross-functional body responsible for overseeing the organization's use and management of artificial intelligence. It establishes policies, accountability mechanisms, and review processes that enable responsible AI adoption while managing legal, security, privacy, ethical, and operational risk."

Also name the authority under which the committee operates, for example the executive leadership team or the risk committee of the board.

2. Scope

Define exactly what falls under the committee's oversight. Be explicit about the range of AI, since risk now comes from far more than in-house models.

- **Built AI:** models and applications developed internally
- **Bought AI:** third-party AI products and services, including AI features embedded in existing SaaS tools
- **Generative AI:** chatbots, coding assistants, and content generation tools
- **Agentic and automated workflows:** systems that take actions or make decisions with limited human input
- **Shadow AI:** unapproved AI tools employees may be using without authorization

> **Example:** A customer support team enables an AI summarization feature that ships inside its ticketing platform. Nobody "bought an AI tool," yet the feature processes customer data through a third-party model. A scope statement that covers embedded AI features brings this under review.

Consider also listing what is out of scope, such as basic rules-based automation, to avoid overreach.

3. Objectives

List four to six measurable goals. Linking each objective to an outcome makes it easier to report progress later.

- Establish and maintain AI policies and standards
- Maintain a complete, current AI inventory
- Apply risk-based review to AI use cases before deployment
- Define and monitor AI risk tolerance
- Provide visibility into unauthorized AI use and address it constructively
- Report significant AI risks to senior leadership

4. Membership and Roles

Membership should bring in the expertise needed to assess AI across departmental boundaries. No single department has all of it. A typical committee has seven roles:

Role	Core responsibilities	Example of their input
Executive Sponsor	Approves governance strategy, sets or approves risk tolerance, provides resources, escalates to the board	Decides the organization will not use AI for automated hiring rejections
AI Governance Chair	Sets the agenda, coordinates risk reviews, maintains documentation, tracks decisions and remediation	Schedules an urgent review when a business unit finds an unregistered AI tool
Legal and Compliance	Assesses regulatory, contractual, IP, and transparency requirements	Flags that a vendor's terms allow training on customer data
Cybersecurity	Evaluates access controls, third-party AI services, prompt injection, and incident response	Requires single sign-on and logging before a chatbot connects to internal systems
Data Governance	Sets data classification, quality, provenance, retention, and privacy requirements	Prohibits uploading restricted data to external AI services

AI/ML or Technology	Model evaluation, testing, monitoring, documentation, and technical controls	Requires a bias and accuracy test before a model goes live
Business or Product	Defines objectives, users, acceptable performance, and human oversight needs	Explains that a tool will be used by frontline staff with no technical background

Add each member's name, title, and term in a roster table at the end of the charter. Where possible, name an alternate for every seat so that meetings are not cancelled when one person is unavailable.

Tip: The chair can come from AI, risk, compliance, technology, or another function. Choose someone with the time, credibility, and cross-functional relationships to follow decisions through, rather than simply the most senior person available.

5. Responsibilities

Document the committee's standing duties. These are the ongoing work, distinct from the objectives above.

- **Policies and standards:** approved tools, prohibited use cases, sensitive-data restrictions, human oversight, third-party requirements, incident reporting, and employee responsibilities. Review them periodically as capabilities and regulations change.
- **AI inventory:** record each system's business owner, purpose, model or provider, data used, risk classification, deployment environment, user population, monitoring needs, and review date.
- **Risk management:** apply a consistent process to identify, assess, treat, and monitor risks such as privacy, security, bias, inaccurate outputs, lack of transparency, IP concerns, and third-party dependencies.
- **Shadow AI governance:** understand why employees turn to unapproved tools and offer safer alternatives rather than only blocking them.
- **Incidents and exceptions:** review AI incidents and approve or deny policy exceptions.
- **Post-deployment monitoring:** confirm that approved systems are monitored for performance, security, and compliance.

Example: Employees are pasting meeting notes into a public chatbot to summarize them. Rather than issuing a blanket ban, the committee approves an enterprise-licensed assistant with data protections, publishes guidance on what may be entered, and adds the tool to the approved registry.

6. Decision Rights and Authority

This is the section that gives the charter teeth. Spell out which decisions the committee can make on its own, which need executive approval, and which are delegated to business owners.

Committee decides	Escalate to Executive Sponsor	Delegated to business owner
Approve or reject medium-risk use cases	High-risk or prohibited-category use cases	Day-to-day operation of approved low-risk tools
Approve standard AI policy updates	Changes to organizational risk tolerance	User training and adoption
Grant time-limited exceptions	Exceptions affecting regulated data	Routine performance checks per the monitoring plan

Also state what happens if the committee cannot reach a decision, for example a simple majority with the chair holding a casting vote, or automatic escalation to the sponsor.

7. Meeting Cadence and Quorum

- **Frequency:** monthly is common at the start, moving to quarterly once processes mature, with ad hoc sessions for urgent reviews.
- **Quorum:** for example, the chair plus at least four members, which must include legal or compliance and cybersecurity.
- **Decision rule:** consensus first, then majority vote if needed.
- **Urgent reviews:** define a fast path, such as a written review by three members within five business days, for time-sensitive requests.

Tip: A slow committee drives people to shadow AI. Publish target turnaround times, for example five business days for low-risk requests, and track them.

8. Escalation Path

Describe when and how issues travel upward. Triggers might include a serious AI incident, a regulatory inquiry, a disagreement the committee cannot resolve, or a proposed use that exceeds the approved risk tolerance.

> **Example:** An AI-assisted customer chatbot begins giving customers incorrect refund information. The chair convenes an urgent review, security and legal assess the exposure, and because customers were affected the chair escalates to the Executive Sponsor within one business day.

9. Reporting and Documentation

Set out what the committee records and what it reports. Good documentation is what turns governance from opinion into evidence.

- **Records to keep:** approvals, assessments, conditions, exceptions, incidents, and monitoring outcomes
- **Regular report to leadership:** number of AI systems by risk tier, use cases reviewed, open exceptions, incidents, and shadow AI findings
- **Retention:** how long records are kept and where they are stored

10. Charter Review and Amendment

AI capabilities and regulations change quickly. Commit to reviewing the charter at least annually, and after any major regulatory change or significant incident. State who may propose amendments and who approves them.

11. Approval Block

Name and title	Signature	Date
[Executive Sponsor]		
[AI Governance Chair]		

Part 2: RACI Matrix

A RACI matrix removes ambiguity about who does what. When responsibilities are unclear, tasks are duplicated, skipped, or argued over at the worst possible moment, usually during an incident.

Legend

Code	Meaning	What it looks like in practice
R	Responsible	Does the work. There may be several.
A	Accountable	Owns the outcome and has final sign-off. Only one per activity.
C	Consulted	Provides input before a decision. Two-way communication.
I	Informed	Told the outcome after the fact. One-way communication.

Sample matrix

The matrix below maps common governance activities to the seven committee roles. Abbreviations: **ES** Executive Sponsor, **Chair** AI Governance Chair, **L&C** Legal and Compliance, **Sec** Cybersecurity, **Data** Data Governance, **Tech** AI/ML or Technology, **Biz** Business or Product.

Activity	ES	Chair	L&C	Sec	Data	Tech	Biz
Approve AI governance policy	**A**	R	C	C	C	C	I
Set AI risk tolerance	**A**	R	C	C	C	C	C

Register a new AI use case	I	I	I	I	I	C	**A/R**
Conduct risk assessment	I	**A**	R	R	R	R	C
Classify risk level	I	**A**	C	C	C	R	C
Approve or reject use case	I	**A/R**	C	C	C	C	I
Implement technical controls	I	I	I	C	C	**A/R**	C
Monitor deployed systems	I	**A**	C	R	R	R	R
Manage AI incidents	I	**A**	R	R	C	R	C
Approve exceptions	**A**	R	C	C	C	C	I
Review shadow AI findings	I	**A**	C	R	C	C	C
Report to executive leadership	**A**	R	C	C	C	C	I

Reading the matrix: two examples

Example: Manage AI incidents. The chair is Accountable, meaning they make sure the response happens and is closed out. Legal, security, and technology are Responsible for the hands-on work: assessing legal exposure, containing the issue, and fixing the system. Data and business are Consulted for context. The sponsor is Informed.

Example: Register a new AI use case. The business owner is both Accountable and Responsible, because the team proposing the AI is best placed to describe what it will do, who will use it, and what data it will touch. Technology is Consulted to help complete the technical details.

Guidance notes

- **One Accountable per activity.** If two people are accountable, nobody is. A/R means one role is both Accountable and Responsible.
- **Avoid over-consulting.** If everyone is Consulted on everything, decisions stall. Consult only those whose expertise changes the outcome.
- **Make sure every row has at least one R.** Someone must actually do the work.
- **Adapt for smaller organizations.** One person may hold several seats, for instance a head of IT covering both security and technology. Combine columns rather than leaving roles unstaffed.
- **Adapt for larger organizations.** Add columns for regional leads, procurement, privacy, or internal audit, and add rows for activities such as third-party AI vendor onboarding.
- **Match RACI to decision rights.** The Accountable role for approvals should agree with the decision-rights table in your charter.

Blank template

Activity	ES	Chair	L&C	Sec	Data	Tech	Biz

Part 3: Meeting Agenda and Minutes

Templates

Consistent meetings produce consistent decisions. These templates keep the committee focused on governance work and leave a clear record of what was decided and why.

1. Standard Meeting Agenda (60 minutes)

Time	Item	Lead	Purpose
0-5	Opening and attendance	Chair	Confirm quorum and note conflicts of interest
5-10	Review of previous actions	Chair	Confirm what is complete, overdue, or blocked
10-30	New AI use case reviews	Business owners	Decide on submitted use cases using the review summary sheet
30-38	High-risk system updates	Tech / Sec	Report on monitoring results and changes to high-risk systems
38-45	Incidents and exceptions	Sec / L&C	Review incidents since the last meeting and any requested exceptions
45-50	Shadow AI and inventory updates	Chair / Data	Share discovery findings and inventory gaps
50-55	Regulatory and policy updates	L&C	Flag upcoming changes that affect policy or use cases
55-60	Decisions and action items	Chair	Read back decisions, owners, and due dates

Tip: Circulate papers at least three business days before the meeting so members arrive prepared. Reserve meeting time for discussion and decisions, not for reading.

2. Use Case Review Summary Sheet

Complete one sheet per AI use case. It gives members the same information in the same format every time, which speeds up review and makes comparisons easier.

Field	Entry
Use case name	[e.g., Customer support reply drafting assistant]
Business owner	[Name, department]
Purpose and expected outcome	[What the system will do and the business result expected]
Users	[Who will use it and how many]
Model or provider	[Internal model, or vendor and product name]
Data processed	[Data types and classification, including any personal or confidential data]
Level of automation	[Advisory only / human review required / fully automated]
Regulatory exposure	[Applicable laws, regulations, contractual obligations]
Risk level	[Low / Medium / High, with rationale]
Key risks identified	[Privacy, security, accuracy, bias, third-party dependency, etc.]
Required controls	[Access restrictions, human review, logging, testing, training]
Decision	[Approved / Approved with conditions / Deferred / Rejected]
Conditions and review date	[Conditions attached and when the system will be reviewed again]

Completed example

Field	Entry

Use case name	Customer support reply drafting assistant
Business owner	Priya Nair, Head of Customer Support
Purpose and expected outcome	Draft first-pass replies to customer tickets, reducing average response time by 30%
Data processed	Customer names, order details, and free-text ticket content (personal data)
Model or provider	Third-party large language model accessed through the ticketing platform
Level of automation	Human review required before any reply is sent
Risk level	Medium: personal data and a third-party provider, but a human reviews every output
Decision	Approved with conditions
Conditions and review date	Vendor confirms customer data is not used for training; reply logs retained for 12 months; review in six months

3. Minutes and Decision Log Template

Minutes should capture decisions and rationale, not a transcript. The decision log matters most, because it is the evidence you will point to in an audit or after an incident.

Meeting record

Field	Entry
Meeting date and time	[Date, time]
Attendees / apologies	[Names, roles; note alternates]
Quorum confirmed	[Yes / No]

Conflicts of interest declared	[None / details]

Decision log

Ref	Decision	Owner	Rationale and conditions	Due	Status
D-001	Approved support reply assistant with conditions	P. Nair	Medium risk; vendor no-training clause required	15 Nov	Open
D-002	Rejected use of public chatbot for contract review	Legal	Confidential data exposure; enterprise tool proposed instead	N/A	Closed

Action items

Action	Owner	Due date	Status
[e.g., Confirm vendor data-training terms in writing]	[Legal]	[Date]	[Open]

Part 4: First 90 Days Checklist

A committee that meets but never operationalizes its processes quickly loses credibility. This checklist follows the lifecycle logic used throughout AI governance (identify, assess, classify, approve, deploy, monitor, and review) and builds the committee's capability in three 30-day phases.

Days 1-30: Foundation

Goal: establish authority, people, and a first view of where AI is being used.

- ☐ Secure an executive sponsor and confirm their time commitment
- ☐ Identify and invite committee members from each function, plus alternates
- ☐ Adapt and adopt the charter, with sponsor sign-off
- ☐ Draft an initial AI inventory from IT records, procurement data, and department surveys
- ☐ Hold the first meeting and agree on meeting cadence

Example: The initial inventory will be incomplete, and that is fine. A common first finding is that teams are using AI features embedded in existing software that nobody thought of as "AI". Record what you find and refine over time.

Days 31-60: Structure

Goal: build the working tools of governance.

- ☐ Finalize the RACI matrix and share it with all stakeholders
- ☐ Define risk classification criteria (impact, data sensitivity, user population, automation, regulatory exposure, business criticality)
- ☐ Launch a use case intake process with a simple registration form
- ☐ Draft or update the AI policy: approved tools, prohibited uses, sensitive-data rules, human oversight
- ☐ Begin shadow AI discovery using SaaS discovery tools, network monitoring, and an employee reporting channel

Example: A three-tier classification works well to start. **Low:** internal productivity tools with no sensitive data. **Medium:** tools processing personal or confidential data with human review. **High:** systems that influence decisions about people, operate with significant automation, or face heavy regulation.

Days 61-90: Operationalize

Goal: run real use cases through the full process and start reporting.

- ☐ Review the first use cases end to end, from registration to documented decision
- ☐ Set the monitoring approach and reporting cadence for approved systems
- ☐ Establish incident and exception handling procedures, including the escalation path
- ☐ Deliver the first report to executive leadership
- ☐ Schedule a 90-day retrospective to review what worked and what to change

Success indicators

Track a few simple measures so you can show progress and spot problems early:

Indicator	What good looks like at 90 days
Inventory coverage	Known AI systems recorded, with owners and risk tiers assigned
Use case reviews completed	First use cases taken through the full workflow
Review turnaround time	Within the published target (for example, five business days for low risk)
Shadow AI visibility	Discovery under way, with approved alternatives identified for the top use cases
Meeting attendance	Quorum met at every meeting
Leadership reporting	First report delivered and acted on

Next Steps and References

Where to go from here

- Build out your AI inventory and risk register so the committee has a reliable basis for decisions. Our AI Inventory and Risk Register Template is a natural companion to this kit.
- Run a pilot review on one real use case to test the charter, RACI, and templates before rolling out more widely.
- Revisit the documents after your 90-day retrospective and adjust them based on what you learned.

Framework references

- **NIST AI Risk Management Framework (AI RMF)** describes governance as a cross-cutting function informing AI risk management throughout the lifecycle, emphasizing policies, accountability structures, defined roles, and AI inventories.
- **NIST Generative AI Profile** identifies risks specific to or intensified by generative AI and recommends oversight, testing, and documentation.
- **ISO/IEC 42001:2023** sets requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS).

Build deeper skills

If you want structured, practical training in AI governance, the GSDC Certified AI Governance Professional certification covers frameworks such as ISO/IEC 42001 and NIST AI RMF, along with AI inventories, risk registers, governance committees, RACI structures, impact assessments, and generative AI governance.

Disclaimer: The templates in this kit are general starting points and do not constitute legal advice. Adapt them to your organization's size, industry, and legal and regulatory context, and consult qualified legal counsel where appropriate.

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