YAKAMICHI INTELLIGENCE CENTRE

PLAYBOOK 01 • RESTORED PUBLICATION

Enterprise AI Governance

A practical operating playbook for governing artificial intelligence across strategy, risk, security, responsible use and continuous assurance.

STRATEGY	GOVERNANCE	RISK	SECURITY	ASSURANCE
Align	Own	Assess	Protect	Monitor

EXECUTIVE PURPOSE

Enable confident AI adoption without separating innovation from accountability. This playbook converts governance principles into decision rights, workflows, controls, templates and a 180-day implementation program.

YAKAMICHI PERSPECTIVE

Governance is the operating system for responsible AI at scale.

Contents & How to Use This Playbook

How to Use It

- Boards and executive teams: use Chapters 1, 3, 9 and 10 to establish oversight, appetite and reporting.
- AI governance leaders: use Chapters 2–8 as the operating backbone for intake, assessment, approval and monitoring.
- Business owners: use the classification, risk and human-oversight tools before deployment and whenever a material change occurs.
- Security, privacy, legal and risk teams: use the appendices as common evidence and decision records rather than maintaining disconnected reviews.

DESIGN PRINCIPLE

Apply governance in proportion to consequence. Low-risk experimentation should remain usable. High-impact AI should receive deeper review, stronger approval, tighter access, more evidence and continuous assurance.

CHAPTER 01

Executive Mandate & Governance Principles

Artificial intelligence is a business capability with strategic, operational, legal, privacy, cyber, ethical and reputational consequences. Governance therefore belongs to enterprise leadership, not to one technical function.

Six Governance Pillars

Strategy Link AI investment and use cases to measurable business outcomes.

Governance Define board oversight, executive sponsorship, ownership and decision rights.

Risk Assess consequence, likelihood, uncertainty and dependencies before deployment.

Security Protect data, models, prompts, APIs, identities and supporting infrastructure.

Responsible AI Embed fairness, transparency, explainability, accountability and meaningful human oversight.

Continuous Assurance Monitor performance, incidents, suppliers, compliance and changing conditions.

Board-Level Decisions

- Approve enterprise AI governance principles and risk appetite.
- Require visibility of material AI systems and high-impact use cases.
- Ensure residual risk has an accountable business owner.
- Require escalation thresholds for material incidents, policy exceptions and unacceptable performance.
- Review whether governance enables safe scaling rather than merely producing compliance artifacts.

OPERATING PRINCIPLE
Automation never removes accountability. Every material AI system must have an identifiable business owner who remains responsible for outcomes.

CHAPTER 02

AI Inventory & Use-Case Classification

Leadership cannot govern systems it cannot see. The inventory should include internally developed models, embedded AI in enterprise software, third-party AI services, automated decision tools, copilots, agents and approved experimental use cases.

Minimum Inventory Record

FIELD	MINIMUM RECORD
Identity	System/use-case name, owner, business function and status
Purpose	Decision or task supported; intended users and beneficiaries
Technology	Model/provider, deployment pattern and integrations
Data	Input/output data categories and sensitivity
Impact	Who or what can be affected by an error or misuse
Controls	Human oversight, security, monitoring and approval status

Risk Tiering

TIER	CHARACTERISTICS	GOVERNANCE
1 — Low	Assistive, reversible, limited sensitive data	Business approval + baseline controls
2 — Moderate	Operational dependency or confidential data	Documented assessment + functional review
3 — High	Consequential decisions, regulated data, material exposure	Cross-functional review + executive owner
4 — Critical	Safety, rights, systemic or enterprise-critical consequence	Executive approval + enhanced assurance

INVENTORY RULE

No material AI system should move into production without a named owner, documented purpose, risk tier and review status.

CHAPTER 03

Operating Model, Roles & Decision Rights

Effective governance separates oversight, approval, execution and assurance while keeping business ownership unmistakable.

ROLE	PRIMARY ACCOUNTABILITY
Board / Committee	Oversight, risk appetite, material issues and strategic alignment
Executive Steering Committee	Policy, prioritization, exceptions and high-risk decisions
AI Governance Office	Standards, inventory, workflow, reporting and coordination
Business Owner	Purpose, outcomes, resources and residual risk
Technology / Data	Architecture, engineering, data quality and technical controls
Security / Privacy / Legal / Risk	Independent domain review and challenge
Internal Audit / Assurance	Periodic independent assessment of design and effectiveness

Decision Rights

- Business owners propose and remain accountable for use cases.
- Governance assigns risk tier and required reviewers.
- Control functions approve within their mandates or record conditions for approval.
- Material exceptions are time-bound and accepted only by authorized risk owners.
- Critical systems and unresolved material risks escalate to executive governance.

AVOID THE COMMITTEE TRAP

A committee can coordinate accountability, but it cannot absorb it. Every decision must still have a named owner and a recorded rationale.

CHAPTER 04

AI Risk Assessment & Approval

Assessment should focus on plausible harm and business consequence rather than generic checklists. The depth of evidence should increase with risk.

DOMAIN	QUESTIONS
Strategic	Does the use case support a defined objective? Could failure impair strategy or reputation?
Operational	What happens when outputs are wrong, unavailable or manipulated? Is there a fallback?
Cyber	Can data, prompts, models, identities, tools or integrations be abused?
Privacy	Is personal or sensitive data necessary, lawful and minimized?
Legal / Compliance	Are decisions, records, disclosures and sector obligations addressed?
Responsible AI	Could bias, opacity, autonomy or weak oversight create material harm?
Third Party	What provider, model, hosting or upstream dependencies can change the risk?

Approval Outcomes

Approve Risk is within appetite and controls are adequate.

Approve with Conditions Deployment may proceed after specified controls or evidence are completed.

Pilot / Restrict Use is limited by population, data, geography, function or duration.

Escalate Residual risk exceeds delegated authority.

Reject / Pause Risk cannot presently be reduced to an acceptable level.

EVIDENCE STANDARD

Record the decision, conditions, accountable owner, residual risk and date for reassessment. Approval should be auditable, not remembered.

CHAPTER 05

Security, Privacy & Data Protection

AI expands the attack surface through prompts, model interfaces, data pipelines, plugins, agents, APIs and third-party services. Security controls should follow the architecture and consequence of the use case.

Core Control Families

CONTROL	EXECUTIVE EXPECTATION
Identity & Access	Least privilege, strong authentication and controlled administrative access
Data Protection	Approved data classes, minimization, encryption and retention controls
Model / Prompt Security	Protect system instructions, secrets, sensitive context and model artifacts
Integration Security	Review APIs, tools, plugins, agents and downstream actions
Logging	Capture security-relevant activity while respecting privacy and retention requirements
Testing	Adversarial testing proportional to risk before and after material change
Incident Response	AI-specific escalation, containment, evidence and recovery procedures

Shadow AI

Unapproved public AI services can create data leakage, contractual and recordkeeping exposure. The response should combine clear policy, approved alternatives, technical controls and practical staff guidance rather than relying on prohibition alone.

SECURITY PRINCIPLE
Treat prompts and model outputs as part of the information environment. Sensitive data does not become less sensitive because it is placed inside an AI workflow.

CHAPTER 06

Responsible AI & Human Oversight

Responsible AI translates values into operating controls. The key question is not whether a system is 'ethical' in the abstract, but whether its design and use preserve appropriate accountability, transparency and human judgment.

PRINCIPLE	PRACTICAL CONTROL
Fairness	Test relevant outcomes and investigate material disparities
Transparency	Disclose AI involvement where appropriate and document intended use
Explainability	Provide explanations suitable to the consequence and audience
Human Oversight	Define who reviews, overrides, escalates or stops the system
Contestability	Provide review or appeal where consequential outcomes affect people
Accountability	Assign business ownership and preserve decision records
Reliability	Define performance thresholds, limitations and fallback procedures

Human Oversight Patterns

- Human-in-the-loop: a person approves before the consequential action occurs.
- Human-on-the-loop: automation proceeds while a person monitors and can intervene.
- Human-in-command: leadership sets boundaries, monitors aggregate outcomes and can suspend the capability.

OVERSIGHT TEST

A human review step is meaningful only if the reviewer has time, information, competence and authority to disagree with the AI.

CHAPTER 07

Third-Party & Foundation Model Governance

Third-party AI can accelerate adoption while introducing opaque dependencies. Procurement must evaluate the service as part of the enterprise control environment, not merely as software functionality.

AREA	DUE-DILIGENCE QUESTIONS
Data Use	Will provider or subprocessors retain, train on or otherwise use organizational data?
Security	What access, encryption, isolation, logging and incident controls apply?
Model Change	Can the underlying model change without notice and alter performance or risk?
Subprocessors	Which hosting, model and service dependencies sit beneath the provider?
Assurance	What testing, certifications, reports or independent evidence are available?
Exit	Can data, configurations and business processes be migrated if the service fails or terms change?

Contractual Requirements

- Data-use and confidentiality restrictions.
- Security and material-incident notification.
- Subprocessor transparency where consequential.
- Change notification for material service or model modifications.
- Audit/assurance rights appropriate to criticality.
- Data return/deletion and transition support at termination.

DEPENDENCY PRINCIPLE

A vendor can operate the model, but the organization retains responsibility for the business use and its consequences.

CHAPTER 08

Lifecycle Monitoring, Incidents & Change

Approval is a point in time. AI risk changes as models, data, prompts, users, integrations, regulations and business context change.

Continuous Monitoring

MONITOR	EXAMPLES
Performance	Accuracy, quality, reliability and task-specific thresholds
Safety / Harm	Material adverse outcomes, prohibited behavior and control failures
Security	Abuse, suspicious access, prompt attacks, tool misuse and data leakage
Fairness	Relevant outcome disparities and drift
Operations	Availability, latency, fallback use and manual overrides
Provider	Model/version changes, outages, terms, incidents and assurance changes

Material Change Triggers

- New model or major version.
- New sensitive data category or population.
- New autonomous action, tool or system integration.
- Expansion into a regulated or higher-consequence decision.
- Material deterioration in performance or control effectiveness.
- Provider, hosting, ownership or contractual change that alters exposure.

Incident Response

AI incidents should connect to existing cyber, privacy, legal, operational and crisis processes. The playbook should define who can suspend the system, preserve evidence, notify stakeholders, communicate externally and authorize restoration.

CHANGE RULE
Material change requires reassessment. A system should not inherit permanent approval from an earlier architecture, model or use case.

CHAPTER 09

Metrics, Board Reporting & Assurance

Executive reporting should emphasize material exposure, control effectiveness and decisions required. Counting experiments or policy acknowledgements alone can create a polished dashboard with little governance value.

METRIC	WHY IT MATTERS
Inventory Coverage	Shows whether leadership can see material AI use
High-Risk Reviews	Shows governance throughput and control coverage
Overdue Remediation	Reveals unresolved exposure
Policy Exceptions	Shows where operating reality diverges from policy
Material Incidents	Connects governance to actual outcomes
Monitoring Coverage	Shows whether approved systems remain under assurance
Third-Party Exposure	Highlights critical external dependencies
Residual Risk	Shows material risks accepted above defined thresholds

Board Reporting Format

- What changed since the last reporting period?
- Which material AI risks are outside appetite or deteriorating?
- Which systems or decisions require board awareness or executive action?
- What significant incidents, exceptions or provider changes occurred?
- What is management doing next, by when, and who owns it?

ASSURANCE MODEL

First line owns the AI use. Second line establishes standards and challenges risk. Third line independently assesses whether governance and controls are designed and operating effectively.

CHAPTER 10

180-Day Enterprise Implementation Program

WINDOW	PRIORITY	DELIVERABLES
0–30 Days	Visibility & ownership	Executive sponsor; material AI inventory; interim acceptable-use rules; priority risks
31–60 Days	Policy & classification	Enterprise AI policy; risk tiers; prohibited/restricted uses; intake workflow
61–90 Days	Operating model	Governance forum; decision rights; cross-functional review; approval records
91–120 Days	Controls & suppliers	Security/privacy controls; third-party standards; human-oversight requirements
121–150 Days	Monitoring & incidents	Metrics; change triggers; AI incident procedures; provider monitoring
151–180 Days	Assurance & board reporting	Executive dashboard; high-risk review; exercise; remediation plan

First Executive Workshop

- Approve the governance objective and risk appetite.
- Select the initial population of material AI systems.
- Agree the four risk tiers and escalation thresholds.
- Name accountable business and governance owners.
- Set the first 90-day deliverables and reporting cadence.

Success at Day 180

The organization should be able to answer: What AI do we materially rely on? Who owns it? What risk tier is it? What evidence supported approval? Which controls are operating? What has changed? Which risks remain outside appetite? What does leadership need to decide?

IMPLEMENTATION PRINCIPLE
Do not wait for a perfect enterprise framework before governing the highest-consequence systems. Establish visibility and ownership first, then mature the machinery around them.

APPENDIX A

AI Use-Case Intake Template

FIELD	RECORD
Business owner	Name / role / function
Use-case title	Short unique identifier
Business objective	Outcome and measurable value
Users / affected parties	Who uses it and who may be affected
AI capability	Model/service and intended function
Data	Input/output categories and sensitivity
Decision / action	What the AI recommends, generates or executes
Human oversight	Reviewer, override and escalation
Dependencies	Provider, integrations and critical upstream services
Proposed risk tier	1 / 2 / 3 / 4 with rationale
Target date	Pilot / production / material change

Use this intake before procurement or production deployment so governance can influence design rather than review a finished decision.

APPENDIX B

AI Risk Assessment Record

ASSESSMENT AREA	RATING / EVIDENCE
Strategic consequence	Low / Moderate / High / Critical + rationale
Operational failure	Failure modes, fallback and recovery
Cybersecurity	Threats, controls, testing and residual exposure
Privacy / data	Lawful basis, minimization, retention and sensitive data
Legal / regulatory	Applicable obligations and counsel review
Responsible AI	Fairness, transparency, explainability and contestability
Human oversight	Pattern, authority, competence and intervention
Third-party risk	Provider assurance, dependencies and exit
Monitoring	Metrics, thresholds, drift and change triggers
Residual risk	Owner, acceptance authority and expiry/review date

DECISION RECORD

Final outcome: Approve / Approve with Conditions / Pilot or Restrict / Escalate / Reject or Pause. Record approvers, conditions, due dates and next reassessment.

APPENDIX C

AI Governance RACI

ACTIVITY	BOARD	EXEC	GOV OFFICE	BUSINESS	CONTROL FUNCTIONS
Risk appetite	A	R	C	C	C
AI policy	I	A	R	C	C
Use-case ownership	I	I	C	A/R	C
Risk classification	I	I	A/R	C	C
High-risk approval	I	A/R	R	R	R/C
Residual risk acceptance	I	A	C	R	C
Monitoring	I	I	C	A/R	C
Independent assurance	A/I	I	C	C	C

R = Responsible, A = Accountable, C = Consulted, I = Informed. Adapt the matrix to organizational structure and delegated authority.

APPENDIX D

Board & Executive Dashboard

INDICATOR	CURRENT	TREND	THRESHOLD / ACTION
Material AI systems inventoried	___%	↑ / → / ↓	Target: ___%
Tier 3–4 reviews complete	___%	↑ / → / ↓	Escalate below ___%
Overdue high-risk remediation	___	↑ / → / ↓	Escalate above ___
Material incidents this period	___	↑ / → / ↓	Board notice criteria: ___
Policy exceptions	___	↑ / → / ↓	Review aging > ___ days
Third-party critical dependencies	___	↑ / → / ↓	Unassured: ___
Systems with active monitoring	___%	↑ / → / ↓	Target: ___%
Residual risks outside appetite	___	↑ / → / ↓	Decision required

Narrative Required

- Top three changes since last period.
- Top three risks requiring executive attention.
- Material provider/model changes.
- Decisions required and accountable owners.
- Progress against 180-day roadmap or annual improvement plan.

APPENDIX E

Executive Governance Checklist

- ☐ AI strategy and governance principles approved
- ☐ Executive sponsor and governance owner assigned
- ☐ Material AI inventory maintained
- ☐ Risk tiers and prohibited/restricted uses defined
- ☐ Cross-functional assessment workflow operating
- ☐ Security, privacy and data controls embedded
- ☐ Responsible AI and human-oversight standards documented
- ☐ Third-party AI due diligence and contract standards operating
- ☐ Monitoring and material-change triggers defined
- ☐ AI incident response integrated with enterprise response
- ☐ Policy exceptions time-bound and visible
- ☐ Residual risk accepted by authorized owners
- ☐ Board/executive reporting established
- ☐ Independent assurance scheduled for material controls

MATURITY TARGET

Move from ad hoc experimentation to repeatable, auditable and adaptive governance. The goal is not maximum process. The goal is sufficient control to scale AI with confidence.

APPENDIX F

AI Governance Maturity Model

LEVEL	CHARACTERISTICS	EXECUTIVE MEANING
1 — Ad Hoc	Unmanaged experimentation; fragmented ownership	Risk is poorly understood
2 — Emerging	Initial policy, reviews and accountable owners	Controls exist but are inconsistent
3 — Managed	Enterprise process, risk tiering and reporting	Decisions are repeatable and auditable
4 — Adaptive	Continuous monitoring, testing and improvement	Governance supports confident scale

Yakamichi Perspective

AI governance is not about restricting innovation. It creates the confidence to deploy AI responsibly at scale while protecting information, supporting compliance and strengthening organizational resilience.

The practical test is simple: leadership should know where material AI is used, who owns the outcome, what evidence supports the decision, what could go wrong, which controls reduce that risk and what signals would cause the organization to intervene.

REQUEST A CONFIDENTIAL CONSULTATION

United States: +1 202 413 3763 | Maryland, USA
Nigeria: +234 811 095 0651 | Abuja, Nigeria

YAKAMICHI • Artificial Intelligence • Cybersecurity • Counter-Surveillance • Corporate Risk Intelligence