OKG INSIGHTS

EXECUTIVE RESOURCE

AI Governance Starter Guide

A Practical Guide for Leaders Building Responsible AI Programs

A practical operating model for visibility, accountability, risk-based decisions, implementation controls, monitoring, escalation, and responsible AI adoption.

One Knowledge. Better Decisions.

Last Reviewed: September 2026

About This Guide

Artificial intelligence is moving into organizations faster than many governance structures are evolving. AI may already be embedded in software, analytics platforms, customer tools, hiring processes, productivity applications, vendor solutions, operational workflows, and employee work practices.

That creates an important leadership challenge: how does an organization take advantage of AI while maintaining appropriate accountability, oversight, and control?

> **The central idea**
>
> Effective AI governance is not simply another policy. It is a practical operating model for understanding where AI is being used, determining how much risk it creates, assigning responsibility, making decisions, monitoring performance, and responding when conditions change.

This guide provides a practical starting point that leaders can adapt to their organization. Governance should reflect organizational size, industry, regulatory environment, risk profile, technology environment, and strategic objectives.

Guide at a Glance

FOUNDATIONS Visibility, accountability, risk management, transparency, oversight, and human judgment.	**LIFECYCLE** Discover, assess, decide, control, deploy, monitor, reassess, and retire.
DECISION MODEL Risk tiers, decision rights, approval conditions, and escalation paths.	**OPERATIONS** Vendor AI, monitoring, incidents, roles, executive reporting, and metrics.
IMPLEMENTATION A practical 90-day starting plan for building a functioning governance process.	**EXECUTIVE TOOLS** A decision checklist leaders can use before approving or materially expanding an AI use case.

1. What AI Governance Should Accomplish

CREATE VISIBILITY Organizations cannot govern AI they do not know exists. Leadership should understand where AI is developed, purchased, piloted, embedded, or used - including through third parties.	**ESTABLISH ACCOUNTABILITY** Every meaningful AI use case should have an accountable owner. Cross-functional review should not result in unclear ownership of the business decision.
MATCH OVERSIGHT TO RISK A productivity assistant and a system influencing employment, healthcare, financial, safety, or regulatory decisions should not receive identical scrutiny.	**ENABLE RESPONSIBLE INNOVATION** Clear boundaries can make adoption easier because teams know what is allowed, what requires review, what information is needed, and who decides.
CREATE REPEATABLE DECISIONS AI decisions should not depend entirely on who happens to be in the room. A repeatable process improves consistency across business units.	**MAINTAIN LIFECYCLE CONTROL** Governance should continue after approval through monitoring, reassessment, incident response, change management, and retirement.

2. Foundations of an AI Governance Program

Accountability

Define who owns the business use case, technology, deployment decision, monitoring, risk acceptance, corrective action, and authority to stop or retire the system.

Oversight

Increase review as potential impact increases. Lower-risk uses may follow streamlined review; higher-risk uses may require multidisciplinary or executive-level review.

Risk Management

Assess privacy, cybersecurity, data governance, reliability, accuracy, bias and fairness, legal and regulatory exposure, financial impact, workforce impact, safety, vendor risk, customer impact, reputation, intellectual property, and explainability as relevant.

Transparency

Leadership should be able to determine what a significant AI system does, why it is used, who owns it, what data it relies on, who may be affected, what decisions it influences, and what major controls are in place.

Human Oversight

Determine where human judgment belongs. Meaningful human review requires appropriate expertise, sufficient information, authority to disagree, time to review, and a way to escalate concerns.

Risk management question

The objective is not to prove that an AI system has zero risk. The objective is to understand the risk, determine how significant it is, identify whether it can be controlled, and define who has authority to accept the remaining risk.

3. The AI Governance Workflow

Governance becomes easier to operationalize when employees and leaders can see how an AI use case moves through the organization.

AI Governance Lifecycle

1	**DISCOVER & REGISTER**	Identify the AI use case, purpose, owner, users, data, and affected stakeholders.
↓		
2	**ASSESS**	Evaluate impact, risk, data, legal/regulatory considerations, technology, and vendor dependency.
↓		
3	**DECIDE**	Approve, approve with conditions, escalate, or decline based on defined decision rights.
↓		
4	**IMPLEMENT CONTROLS**	Complete required safeguards, testing, documentation, training, and vendor conditions.
↓		
5	**DEPLOY**	Release the approved capability into operation after required conditions are met.
↓		
6	**MONITOR & REPORT**	Track performance, incidents, changes, complaints, vendor changes, and emerging risks.
↓		
7	**REASSESS / MODIFY /**	Return material changes to review and formally change, suspend, or retire the

SUSPEND / RETIRE	capability when appropriate.

Decision trigger

If the use case, model, vendor, data, affected population, regulation, or operating environment changes materially, return the use case to ASSESS before continuing.

Stage 1 - Discover & Register

Capture the business purpose, accountable owner, users, technology or vendor, data involved, affected individuals or groups, expected benefit, decisions influenced, and implementation status. The output is an AI use-case record that becomes part of the AI inventory.

Stage 2 - Assess

Evaluate the use case in proportion to potential impact. Consider who may be affected, what decisions are influenced, what data is involved, whether failure could materially disrupt the organization, what legal or regulatory requirements apply, and how reliable, secure, explainable, and controllable the system is.

4. A Simple AI Risk-Tier Model

Risk tiers help determine how much oversight an AI use case should receive.

Risk Level	General Characteristics	Governance Approach
Tier 1 - Lower Risk	Limited organizational or stakeholder impact; primarily administrative or productivity use.	Streamlined review and standard controls.
Tier 2 - Moderate Risk	Influences business processes, customer interactions, operational recommendations, or sensitive information.	Formal assessment and designated functional review.
Tier 3 - Higher Risk	Influences consequential decisions, regulated activity, employment, healthcare, financial outcomes, safety, or material organizational exposure.	Enhanced assessment, multidisciplinary review, senior decision authority, stronger documentation, and ongoing monitoring.

Core principle

As potential impact increases, governance scrutiny should increase.

5. The Governance Decision Gate

After assessment, the use case should reach a defined decision point. Four basic outcomes create clarity and avoid indefinite review cycles.

APPROVE Proceed under established requirements.	**APPROVE WITH CONDITIONS** Proceed after required safeguards, testing, documentation, or other conditions are completed.
ESCALATE Refer the use case to a higher decision authority because of risk, novelty, uncertainty, or strategic significance.	**DECLINE** Do not proceed because risk exceeds tolerance, controls are inadequate, or the business value does not justify exposure.

Governance principle

The function assessing risk does not always need to be the authority accepting risk. Define who may approve each risk tier before difficult decisions arise.

6. Implementation and Control

Approval should translate into operational requirements. Controls should reflect the specific use case and the level of risk.

ACCESS CONTROLS Restrict who can use the AI capability and what privileges they have.	**DATA CONTROLS** Define what information may and may not be entered into or processed by the system.
HUMAN REVIEW Require validation before specified outputs are acted upon when appropriate.	**TESTING** Evaluate system performance, limitations, and control effectiveness before production use.
DOCUMENTATION Maintain records of purpose, ownership, assessment, approval, limitations, controls, and monitoring.	**TRAINING** Ensure users understand permitted use, limitations, escalation requirements, and prohibited activities.
VENDOR REQUIREMENTS Address contractual, security, privacy, performance, notification, audit, and change-management expectations.	**PERFORMANCE THRESHOLDS** Define what acceptable performance looks like and when intervention is required.

7. Third-Party and Vendor AI

AI can enter an organization through software updates, cloud platforms, productivity applications, HR systems, customer platforms, cybersecurity tools, analytics, professional services, medical or operational technology, and embedded assistants. The organization may not have built the model, but it can still own the consequences of using it.

WHAT AI IS THE VENDOR PROVIDING? Understand the actual capability rather than relying only on product terminology.	**WHAT ORGANIZATIONAL DATA WILL THE VENDOR RECEIVE?** Determine whether sensitive, confidential, regulated, proprietary, employee, or customer information is involved.
HOW CAN THE VENDOR USE OUR DATA? Determine whether information may be retained, reused, shared, or used to train models.	**WHAT HAPPENS WHEN THE MODEL CHANGES?** Define whether material changes require notification, additional testing, or reassessment.
CAN AI FUNCTIONALITY BE DISABLED? Not every embedded AI feature needs to be activated.	**WHAT HAPPENS TO DATA WHEN THE RELATIONSHIP ENDS?** Understand retention, portability, and deletion requirements.

Accountability reminder

Outsourcing technology does not necessarily outsource accountability.

8. Monitoring: Governance After Approval

Technology, data, vendors, business processes, laws, and patterns of use change. Approved AI systems therefore need ongoing monitoring and defined reassessment triggers.

MONITOR PERFORMANCE Accuracy, reliability, performance thresholds, significant errors, and unexpected outcomes.	**MONITOR STAKEHOLDER SIGNALS** Complaints, fairness concerns, inappropriate outputs, or changes in affected populations.
MONITOR TECHNOLOGY & DATA Model changes, data changes, integrations, security issues, and changes in intended use.	**MONITOR EXTERNAL CONDITIONS** Regulatory developments, vendor changes, contractual changes, and new organizational requirements.

Common Reassessment Triggers

- A new use case or material expansion
- A material model or vendor change
- A new data source
- Use with a new population
- A significant incident
- A regulatory or legal change
- Material performance deterioration
- Expansion into a consequential decision process

9. AI Incident and Escalation Workflow

1	**ISSUE IDENTIFIED**	Potential AI-related problem, control failure, unauthorized use, harmful output, or other significant concern.
2	**CONTAIN**	Stop or limit immediate exposure when appropriate.
3	**ESCALATE**	Notify the appropriate business, technology, risk, privacy, legal, security, compliance, or leadership functions.
4	**ASSESS IMPACT**	Determine what happened, who or what was affected, and the potential severity.
5	**CORRECT & DOCUMENT**	Implement corrective action and maintain a record of the issue, response, findings, and decisions.
6	**DECIDE**	Continue, modify, add controls, suspend, retire, or return the use case to governance review.

Examples of Potential AI Incidents

- Sensitive information disclosed improperly
- Significant inaccurate output
- Unauthorized AI use
- Unexpected discriminatory or biased outcomes
- Material cybersecurity exposure
- AI operating outside its approved purpose
- Failure of required human oversight
- Unexpected vendor or model behavior
- Harmful or inappropriate customer-facing output
- Significant business decisions based on unreliable results

Define thresholds

Not every AI error is necessarily an enterprise incident. Define thresholds for what requires escalation and who must be notified.

10. AI Governance Roles

AI governance is usually cross-functional, but involving more functions does not automatically create better governance. Roles should be explicit and decision rights should remain clear.

Function	Typical Governance Responsibility
Executive Leadership	Establish strategic direction and organizational risk tolerance.
AI Governance Committee / Council	Oversee significant AI risks and cross-functional decisions.
Business Owner	Own the use case, business outcome, and operational accountability.
Technology / AI Team	Evaluate technology, architecture, integration, and technical performance.
Data Governance	Address data quality, lineage, access, and appropriate use.
Cybersecurity	Assess security exposure and controls.
Privacy	Evaluate collection, processing, sharing, retention, and protection of information.
Legal / Compliance	Evaluate legal, contractual, and regulatory obligations.
Enterprise Risk	Support risk assessment, classification, treatment, and escalation.
Human Resources	Review workforce-related implications when applicable.
Procurement / Vendor Management	Address third-party AI and contractual requirements.
Internal Audit / Assurance	Provide independent assurance where appropriate.

11. What Executives Should See

Executives generally do not need to review every AI interaction. They do need visibility into meaningful exposure, unresolved decisions, and assurance signals.

AI PORTFOLIO	GOVERNANCE	ASSURANCE
• Total active AI use cases • New use cases submitted • Use cases under review • Higher-risk AI applications • Third-party AI applications • AI use cases by business unit	• Use cases awaiting decision • Conditional approvals • Overdue reviews • Open high-risk issues • Applications requiring executive review	• AI incidents • Overdue reassessments • Material vendor changes • Systems outside performance thresholds • Open corrective actions

Useful Governance Metrics

INVENTORY COVERAGE Percentage of identified AI systems entered into the AI inventory.	**PRE-DEPLOYMENT REVIEW** Percentage of applicable AI use cases reviewed before implementation.
HIGH-RISK EXPOSURE Number and percentage of AI applications classified as higher risk.	**REVIEW TIMELINESS** Average time from governance intake to decision.
OUTSTANDING CONDITIONS Number of approved AI applications with unresolved governance conditions.	**REASSESSMENT COMPLIANCE** Percentage of required reassessments completed on time.

THIRD-PARTY EXPOSURE	INCIDENTS & CORRECTIVE ACTIONS
Number of AI-enabled third-party applications.	Number and severity of significant incidents, plus open and overdue corrective actions.

Metric design

If a measure does not help leadership understand risk, performance, or required action, it may not belong on the executive dashboard.

12. A Practical 90-Day Starting Plan

Organizations do not need to build a mature AI governance program in one step. A focused first 90 days can establish the operating foundation.

DAYS 1-30	DAYS 31-60	DAYS 61-90
ESTABLISH VISIBILITY	**BUILD THE PROCESS**	**OPERATIONALIZE**
• Identify executive sponsorship • Identify initial governance owners • Begin the AI inventory • Identify immediate higher-risk uses • Assess existing policies	• Define risk categories • Develop an AI intake process • Define decision authority • Establish review criteria • Define baseline controls • Establish initial AI policy	• Begin formal reviews • Test escalation paths • Establish monitoring • Create executive reporting • Establish incident procedures • Review and refine the process
Outcome: The organization understands its initial AI environment and who is responsible for moving governance forward.	**Outcome: The organization has a repeatable process for evaluating whether AI use cases should proceed.**	**Outcome: The organization has moved beyond policy and created a functioning AI governance operating model.**

13. Executive AI Governance Checklist

Business Purpose

- ☐ What business problem are we solving?
- ☐ Is AI appropriate for this problem?
- ☐ What measurable benefit do we expect?

Ownership

- ☐ Who is the accountable business owner?
- ☐ Who owns technical implementation?
- ☐ Who has authority to approve the risk?

Data

- ☐ What data will the AI use?
- ☐ Does it include sensitive, confidential, proprietary, personal, or regulated information?
- ☐ Is the data appropriate for the intended purpose?

People and Impact

- ☐ Who could be affected by the AI system?
- ☐ Could it influence a consequential decision?
- ☐ Could errors create material harm?

Risk

- ☐ What are the most significant risks?
- ☐ Have appropriate functions reviewed those risks?
- ☐ Are controls proportionate to potential impact?

Human Oversight

- ☐ Where is human judgment required?
- ☐ Can a human meaningfully challenge or override an AI outcome?
- ☐ Who can stop the process if necessary?

Vendor

- ☐ Is a third party involved?
- ☐ How may the vendor use organizational data?
- ☐ What happens if the vendor or underlying model changes?
- ☐ Are contractual protections appropriate?

Monitoring and Lifecycle

- ☐ How will performance be monitored?
- ☐ What would indicate the AI is no longer operating as intended?
- ☐ What happens if something goes wrong?
- ☐ What would trigger reassessment?
- ☐ How can the capability be suspended or retired?

14. Governance Should Enable Responsible Adoption

One risk of AI governance is creating a process so complicated that employees begin working around it. Governance should be rigorous where rigor is necessary, but it should also be understandable and proportionate.

- Where do I go with an AI idea?
- What requires approval?
- What information do I need?
- Who can make the decision?
- What can I do without escalation?
- What conditions trigger additional review?

A low-impact use case should not necessarily follow the same path as an AI system influencing consequential decisions. Risk-based governance directs organizational attention toward the areas where the consequences are greatest.

15. Framework Alignment

Organizations do not have to build AI governance from a blank page. Recognized frameworks and standards can inform program development while organizations also account for applicable laws, regulations, contracts, industry requirements, and internal policies.

NIST AI Risk Management Framework

The U.S. National Institute of Standards and Technology developed the AI Risk Management Framework (AI RMF) as a voluntary, flexible resource to help organizations manage AI-related risks. Its core organizes risk management around Govern, Map, Measure, and Manage. As of September 2026, NIST states that AI RMF 1.0 is being revised.

NIST - AI Risk Management Framework

ISO/IEC 42001:2023

ISO/IEC 42001:2023 specifies requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS). It provides a management-system approach to AI-related risks and opportunities and emphasizes continual improvement.

ISO - ISO/IEC 42001:2023

> **Practical alignment**
>
> Use external frameworks to strengthen the operating model - not to create disconnected compliance processes that sit outside how AI decisions are actually made.

16. The Starting Point

AI governance does not begin with a committee, a technology platform, or a lengthy policy. It begins with visibility and accountability.

KNOW WHERE AI IS USED Create enough visibility to understand significant AI use, embedded vendor capabilities, pilots, and higher-impact employee uses.	**UNDERSTAND IMPACT** Determine what decisions the system influences, what data it uses, who may be affected, and what could go wrong.
ASSIGN ACCOUNTABILITY Name an accountable business owner and define who can approve, escalate, stop, or retire the use case.	**MAKE DELIBERATE DECISIONS** Use risk tiers and decision rights so scrutiny increases as potential impact increases.
MONITOR WHAT HAPPENS NEXT Track performance, incidents, changes, complaints, vendor developments, and reassessment triggers.	**KEEP GOVERNANCE USABLE** Make the process clear enough that teams know how to comply without creating unnecessary friction.

One principle to remember

AI governance should make accountability clearer - not make innovation impossible.

About OKG Insights

OKG Insights provides practical perspectives on artificial intelligence, governance, risk, business strategy, operations, and workforce readiness to help organizations navigate increasingly complex environments.

One Knowledge. Better Decisions.

This guide is provided for general informational and educational purposes. Organizations should evaluate applicable legal, regulatory, contractual, industry, and organizational requirements when developing AI governance programs.