

The OSAIL AI Governance Framework

Six principles, five functions, and a 365-day path to AI your people can trust. A practical guide to establishing ethical AI oversight and compliance structures within your organization.

Every AI decision is a leadership decision in disguise.

A system screens a candidate, flags a claim, or ranks a request, but a person chose to let it. This framework exists so that person is never "nobody in particular."

Most organizations did not decide to adopt AI. It arrived. It came through a vendor update, a free trial, or an employee who found a faster way to write a report. By the time leaders asked, "What's our AI policy?", the honest answer was usually, "Whatever each person decided on their own."

That is not a technology problem. It is a governance gap, and governance gaps close the same way they always have: clear principles, named owners, practical routines, and leaders who model the behavior they expect.

The OSAIL AI Governance Framework gives you that structure in three layers. **Six principles** define what you stand for. **Five functions** turn those principles into daily operations. **A 365-day roadmap** shows you what to do first, next, and after that.

How to use this guide. Read Sections 1 and 2 with your leadership team; they take about twenty minutes. Assign Sections 3 and 4 to the person who will own AI governance day to day. Use Section 5 to plan your first year, and complete the self-assessment in Section 7 before you start, then again at 180 and 365 days to measure progress.

CONTENTS

01	Why governance, why now	3
02	The framework at a glance	4
03	The six guiding principles	5
04	The five core functions	8
05	The 365-day readiness roadmap	11
06	Roles and risk tiers	12
07	Governance readiness self-assessment	13
	Sources and next steps	14

Why governance, why now

AI is already inside your organization. The only open question is whether anyone is steering.

75%

of knowledge workers reported using generative AI at work in Microsoft and LinkedIn's 2024 global survey.

78%

of those AI users said they were bringing their own AI tools to work, often without formal approval or guidance.

Those two numbers describe the gap in plain terms. People are moving faster than policy. That speed is not the enemy; it is energy looking for direction. Without guidance, though, well-meaning employees paste client data into public tools, managers approve AI-drafted decisions they can't explain, and no one knows which systems the organization actually depends on.

The rules are moving. Your principles shouldn't.

As of September 2026, the regulatory picture is in active motion:

- In the European Union, lawmakers reached a provisional agreement in 2026 to postpone key high-risk obligations under the EU AI Act. The obligations were deferred, not cancelled.
- In the United States, a December 2025 executive order set out to challenge state AI laws, and the White House followed with a national policy framework in 2026. Colorado replaced its first AI law in 2026 after enforcement was blocked in federal court.
- Closer to home, New York City has required bias audits and candidate notices for automated employment decision tools since 2023.


Any leader who builds governance around a single law will rebuild it every year. The better approach is to anchor governance in principles and routines that satisfy the intent behind all of these rules, then adjust the details as specific requirements settle.

This framework aligns with the two most widely used reference points: the U.S. National Institute of Standards and Technology's AI Risk Management Framework (NIST AI RMF 1.0) and the international management system standard for AI, ISO/IEC 42001. If you later pursue formal certification or face an external audit, the structures you build here will carry forward.

A story worth remembering: in 2018, Amazon scrapped an experimental recruiting tool after finding it downgraded résumés that included the word "women's." It had learned from a decade of hiring data that favored men. The model did what it was trained to do. What was missing was a leader asking whose history it was learning from.

The framework at a glance

Think of it as a building. Principles are the foundation, functions are the structure, and the roadmap is the construction schedule.



How the layers work together

Layer	The question it answers	What you produce
Principles	What do we stand for, and what will we refuse to do?	A board-approved statement of AI principles and an AI acceptable use policy
Functions	Who does what, and how does work flow day to day?	A governance council, named system owners, an AI inventory, an intake process, and review routines
Roadmap	What do we do first, and how do we know we're ready for the next step?	A sequenced first-year plan with clear readiness checks at each milestone

Start small, on purpose. You do not need every structure in place before anyone uses AI. You need clarity on the highest-risk uses first, and a visible commitment that governance will grow alongside adoption. Governance that arrives all at once usually arrives as paperwork, and people route around paperwork.

The six guiding principles

Each principle includes what it means, what it looks like in practice, the warning sign that it's slipping, and one question your leadership team should be able to answer.

1 Transparency

If you can't explain it, you shouldn't deploy it.

IN PRACTICE

- Tell people when AI plays a role in decisions that affect them, including hiring, performance, pricing, and service.
- Keep a plain-language summary for every significant AI system: what it does, what data it uses, and where its limits are.
- Require vendors to explain how their tools reach results before you sign.

WARNING SIGN

The most common answer to "How does this tool decide?" is "The vendor would know."

Leadership question: Could a frontline manager explain our most important AI system to an affected employee in two minutes?

2 Accountability

Every AI outcome has an owner.

IN PRACTICE

- Assign one named system owner, a person and not a department, before any AI system goes live.
- Give that owner the authority to pause the system and the budget to fix it.
- Record every significant AI decision: who approved it, when, and on what evidence.

WARNING SIGN

When something goes wrong, the conversation starts with "Whose system is this?"

Leadership question: For each AI system we rely on, can we name the one person who answers for it?

3 Equity

AI must work for everyone it touches.

IN PRACTICE

- Test results across different groups of people before launch and at regular intervals after.
- Treat historical data with healthy suspicion. It records yesterday's decisions, including the unfair ones.
- Include the people most affected by a system in its design and review, so they are seen, counted, and treated fairly.

WARNING SIGN

The system performs well "on average," and no one has checked who sits below the average.

Leadership question: Who could this system leave out, and how would we find out?

4 Integrity

Values don't pause for efficiency.

IN PRACTICE

- Check every proposed AI use against your stated values, not only your budget.
- Require a formal decision before a pilot becomes permanent.
- Be honest in public about what your AI does and does not do. Never overstate it to customers, funders, or employees.

WARNING SIGN

A "temporary" pilot has quietly been running for a year with no review.

Leadership question: Would we be comfortable if our AI practices appeared on the front page tomorrow?

5 Confidentiality

Data is borrowed trust.

IN PRACTICE

- Publish clear rules on what staff may and may not enter into AI tools, with examples.
- Keep a list of approved tools, and make the approved path easier than the unapproved one.
- Require a data agreement before any client, employee, or health information enters an AI system.

WARNING SIGN

Staff use personal AI accounts for work because the approved option is slower or unclear.

Leadership question: Do we know which AI tools our people actually use, not just the ones we bought?

6 Human-in-the-Loop

AI advises. People decide.

IN PRACTICE

- Define which decisions always require a human, especially those affecting employment, access to services, money, or safety.
- Give reviewers the time, training, and authority to disagree with the system.
- Track how often reviewers override AI output. A rate near zero deserves a closer look.

WARNING SIGN

Reviewers approve nearly everything. A human who approves everything isn't oversight. It's decoration.

Leadership question: Where could a person stop our AI today, and do they know they're allowed to?

Put the principles to work in one meeting. Pick your single most important AI use. Walk through the six leadership questions above as a team. Any question you can't answer confidently is your first governance priority.

The five core functions

Principles tell people what matters. Functions make sure it happens on an ordinary Tuesday, not just in the policy binder.

01

Governance and Oversight

Define who decides, what rules apply, and how the organization sees its AI footprint as a whole.

KEY STRUCTURES

- **AI Governance Council.** A small, cross-functional group that meets monthly: an executive sponsor, plus leaders from legal or compliance, HR, IT and security, data, and operations. Include at least one voice from the front line. Governance designed only at the top tends to miss how work really gets done.
- **AI acceptable use policy.** Two to four pages in plain language, covering approved tools, prohibited uses, data rules, and how to ask for help.
- **AI inventory.** A living list of every AI system in use, its owner, its purpose, its data, and its risk tier (see Section 6).

Output: A chartered council, an approved policy, and a first inventory.

02

Operational Integration

Build governance into the way work already flows, so doing the right thing is also the easy thing.

KEY STRUCTURES

- **AI intake process.** A short request form for any new AI use, routed by risk tier. Low-risk uses get approved in days, not months.
- **Review gates.** Checkpoints before purchase, before launch, and before any pilot becomes permanent.
- **Vendor due diligence.** A standard set of questions for every AI vendor: How does the tool reach results? What data is it trained on? Will our data train it? How do you test for bias? What happens when it fails?

Output: An intake form, a tiered approval path, and a vendor question set used on every purchase.

03

Capability Building

Governance only works when people understand it. Learning should match each role's real responsibilities.

Audience	What they need to know
Board and executives	Strategic risk and opportunity, their oversight duties, and the right questions to ask management
People managers	How to review AI-assisted work, when to escalate, and how to talk with their teams about AI and job security honestly
Practitioners and system owners	Testing, documentation, monitoring, and incident response for the systems they own
All staff	The acceptable use policy, data rules, approved tools, and how to raise a concern without fear

Output: Role-based learning paths, with completion tracked for high-risk roles.

04

Client Advisory

Bring expert guidance to the point of decision, whether the "client" is an internal team asking for help or the organization facing a complex choice.

KEY STRUCTURES

- **A clear advice channel.** One named place, whether a person, an inbox, or regular office hours, where any employee can ask, "Is this AI use okay?" and get an answer within days.
- **Regulatory watch.** A quarterly briefing for the council on legal and standards changes that affect your sector and locations.
- **Independent perspective.** For high-stakes decisions, bring in advisors who have no stake in the tool being approved.

Output: A staffed advice channel and a standing quarterly regulatory briefing.

05

Continuous Review

AI governance is not a one-time exercise. Systems drift, data changes, and rules evolve.

KEY STRUCTURES

- **Monitoring.** Each system owner tracks a few measures: accuracy, fairness across groups, human override rates, and complaints.
- **Incident response.** A simple playbook for when AI causes harm: pause, assess, notify, fix, and learn. Practice it once before you need it.
- **Annual review.** Every system in the inventory is reviewed at least yearly and either renewed, improved, or retired.

Output: A monitoring dashboard, a tested incident playbook, and an annual review calendar.

If you only do three things this quarter

1. **Build the inventory.** You cannot govern what you cannot see. Ask every manager which AI tools their teams use, and make it safe to answer honestly.
2. **Name the owners.** Put one person's name next to every system that touches hiring, performance, customers, or money.
3. **Publish the rules.** A two-page acceptable use policy people actually read beats a forty-page policy nobody opens.

The 365-day readiness roadmap

From first assessment to AI governance in full operation. Each milestone ends with a readiness check, so you advance on evidence, not the calendar.

7
DAYS

Initial Assessment

Find out where you really stand. Complete the self-assessment in Section 7 with your leadership team. Survey managers on which AI tools their teams use, and assure them the goal is visibility, not blame. Name an executive sponsor.

Ready to move on when: you have a first list of AI uses, a completed self-assessment, and a named sponsor.

30
DAYS

Framework Development

Build the essentials. Charter the AI Governance Council. Adopt your statement of principles. Draft and approve the acceptable use policy. Assign a risk tier and a named owner to every system on your list.

Ready to move on when: the policy is approved and published, and every high-risk system has an owner.

90
DAYS

Team Training

Build shared understanding. Roll out role-based learning, starting with executives and people managers. Open the advice channel. Launch the intake process and approved tools list.

Ready to move on when: all managers have completed training and new AI requests are flowing through intake.

180
DAYS

Rollout

Put governance to work across the organization. Apply review gates to every new purchase and launch. Begin monitoring high-risk systems. Run a practice incident to test your playbook. Repeat the self-assessment to measure progress.

Ready to move on when: high-risk systems report monitoring results to the council and your self-assessment score has risen.

365
DAYS

Continuous Excellence

Make governance permanent. Complete the first annual review of every system in the inventory. Refresh the policy based on what you learned. Report results to the board, and celebrate the teams that raised concerns early.

You've arrived when: governance runs as a routine, not a project, and people bring AI questions forward without being asked.

Roles and risk tiers

Who does what

Role	Core responsibilities
Board	Approves AI principles and risk appetite. Receives an annual governance report.
Executive sponsor	Champions governance, secures resources, and chairs or appoints the chair of the council.
AI Governance Council	Sets policy, approves high-risk uses, reviews the inventory, and oversees incidents.
System owner	Answers for one AI system: its documentation, monitoring, reviews, and fixes. Holds authority to pause it.
People managers	Guide AI use within their teams, review AI-assisted work, and escalate concerns.
All employees	Follow the acceptable use policy and speak up when something looks wrong.

Match oversight to risk

Not every AI use deserves the same scrutiny. A tool that drafts meeting notes is not a tool that screens job applicants. Tiering keeps governance fast where it can be and careful where it must be.

Tier	Typical examples	Approval	Review
Low	Drafting internal documents, summarizing public information, brainstorming	Approved tools list; no individual approval needed	Annual policy check
Medium	Customer-facing content, internal analytics, productivity tools that touch internal data	System owner and manager sign-off through intake	Annual system review
High	Hiring, performance, lending, pricing, eligibility, health, safety, or any decision that significantly affects a person	Council approval, testing for fairness, and documented human oversight	Ongoing monitoring plus a review at least every six months

When in doubt, tier up. Moving a system to a lower tier after review costs a few days. Discovering a high-risk system was treated as low risk can cost your reputation.

Governance readiness self-assessment

Score each statement from 1 to 4. Complete it as a leadership team, and be honest. The value is in the conversation, not the number.

1 Not in place · 2 Started · 3 Mostly in place · 4 Fully in place and working

#	Statement	1	2	3	4
1	We have a current list of the AI tools and systems used across the organization.	☐	☐	☐	☐
2	Every significant AI system has one named owner with authority to pause it.	☐	☐	☐	☐
3	We have a written, approved AI acceptable use policy that staff know about.	☐	☐	☐	☐
4	Staff know which AI tools are approved and what data they may enter.	☐	☐	☐	☐
5	People are told when AI plays a role in decisions that affect them.	☐	☐	☐	☐
6	We test high-risk AI systems for fair outcomes across different groups.	☐	☐	☐	☐
7	High-stakes decisions always include a person with real authority to override AI.	☐	☐	☐	☐
8	We ask every AI vendor a standard set of governance questions before purchase.	☐	☐	☐	☐
9	Leaders and managers have received AI governance training suited to their roles.	☐	☐	☐	☐
10	Employees can raise AI concerns easily and without fear of consequences.	☐	☐	☐	☐
11	We have a plan for what to do when an AI system causes harm.	☐	☐	☐	☐
12	Our AI systems are reviewed at least once a year and retired when they no longer serve us.	☐	☐	☐	☐

Total score	Stage	What it means
12 to 23	Exposed	AI is in use without meaningful guardrails. Start with the 7-day and 30-day milestones now.
24 to 35	Emerging	Foundations exist but are uneven. Focus on ownership, training, and your high-risk systems.
36 to 43	Established	Governance is working. Strengthen monitoring, fairness testing, and incident readiness.
44 to 48	Leading	Governance is part of how you operate. Share what you've learned, and keep reviewing.

SOURCES

References

Dastin, J. "Amazon scraps secret AI recruiting tool that showed bias against women." Reuters, October 10, 2018.

Microsoft and LinkedIn. "AI at Work Is Here. Now Comes the Hard Part." 2024 Work Trend Index. May 2024.

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). January 2023.

International Organization for Standardization. ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system. December 2023.

Sidley Austin. "EU Lawmakers Reach Provisional Agreement to Delay Key EU AI Act Obligations." Data Matters, June 22, 2026.

The White House. "Ensuring a National Policy Framework for Artificial Intelligence." Executive order, December 2025.

McDermott Will & Schulte. "Colorado AI law in flux: Comprehensive replacement bill signed after federal court blocks predecessor's enforcement." 2026.

NYC Department of Consumer and Worker Protection. Automated Employment Decision Tools (Local Law 144). Enforced since July 2023.

Governance is a leadership practice, not a document.

This framework gives you the structure. Putting it into action inside a real organization, with its history, its politics, and its people, is where most efforts stall. OSAIL works alongside leadership teams to assess readiness, build governance that fits, and prepare people to lead with confidence.

Questions, or ready to talk about your organization?

info@odsynergisticsconsultingllc.com

© 2026 OD Synergistics Consulting LLC. All rights reserved. OSAIL is the Office of Strategic AI Integration and Leadership, part of OD Synergistics Consulting LLC. You may download, print, and share this framework within your organization for internal use, with attribution intact. You may not sell, license, rebrand, or repackage it, or include it in paid products or client deliverables, without written permission. This guide is provided for general educational purposes and is not legal, regulatory, or compliance advice. Regulatory information reflects publicly reported developments as of September 2026 and may change.