Defensible

The Defensible AI Adoption Readiness Framework

Built for non-profits

A Practical Guide to Enabling AI
Without Compromising Control

WHAT'S INSIDE

01 AI Is Already Inside Your Organization

Most organizations didn't decide to adopt AI. It arrived on its own.

Staff started using ChatGPT to draft emails. Someone connected Copilot to their Microsoft 365 account. A vendor quietly embedded AI into a platform the organization had been using for years. By the time leadership began discussing an AI strategy, the tools were already running.

This is the reality of Shadow AI and it's where every honest conversation about AI governance has to begin.

> *Banning AI doesn't make your organization safer. It just makes the risk invisible.*

THE PATTERN MOST ORGANIZATIONS FOLLOWED

When AI tools became widely available, organizations broadly went through four stages:

Discovery. Leadership realized staff were already using AI tools informally on personal devices, through browser extensions, or via embedded vendor features.

Restriction. The instinct was to ban it. Many organizations issued blanket prohibitions, concerned about data exposure and lack of control.

Reality check. Bans didn't hold. Staff continued using tools informally, often in ways that were harder to monitor than sanctioned use. The risk didn't go away it went underground.

Policy and governance. The organizations that got it right stopped trying to prevent AI use and started shaping it. They replaced prohibition with structure.

Many organizations are still working through this transition. The goal of this framework is to help complete it.

WHAT THIS MEANS PRACTICALLY

Culture shapes AI adoption more than technology does. If staff don't understand what's permitted and why, they'll make their own decisions. Policy works. Prohibition doesn't.

02 Where Risk Actually Builds

AI risk isn't theoretical. It builds in specific places, through specific behaviors. Understanding where exposure originates is the foundation of a credible governance program.

There are four categories to address.

DATA RISKS

The most immediate exposure comes from what people put into AI tools. Staff routinely paste sensitive information donor lists, financial reports, board communications, grant documentation into commercial AI platforms without considering where that data goes or how it's stored.

Most commercial AI tools retain inputs for model training by default. Without explicit configuration or enterprise agreements, your organization's sensitive information may be used in ways you didn't anticipate and can't reverse.

Before enabling AI broadly, every organization should be able to answer: What information are we entering? Where is it stored? Who else could access it? Does it leave our control?

TECHNICAL RISKS

AI tools don't operate in isolation. They connect to the environments where your data already lives Microsoft 365, Google Workspace, CRM platforms, fundraising systems, volunteer databases. Each integration is a potential exposure point.

The most common technical risks are unsecured integrations, weak authentication, shadow AI applications running without IT oversight, and overpermissioned accounts that give AI tools access to more data than necessary.

COMPLIANCE RISKS

Depending on your organization's work, AI use may intersect with HIPAA, FERPA, state privacy laws, grant reporting requirements, or donor confidentiality agreements. These aren't abstract concerns.

Entering protected data into a public AI tool, publishing AI-generated content without human review, or sharing restricted information with third-party vendors can create real compliance exposure often without any awareness that a line was crossed.

PEOPLE AND USER RISKS

The most persistent risk is over-reliance: staff trusting AI outputs without verification, publishing AI-generated content without review, or assuming that automation equals accuracy.

> *AI can assist, but it should not decide.*

The organizations that use AI well maintain a clear principle: humans remain responsible for every output. AI accelerates the work. It doesn't replace the judgment required to do it right.

03 Readiness Starts With Governance

Before your organization enables AI broadly, there's a foundational question that needs an honest answer: Do you know what data AI tools can already see?

Most organizations don't. That gap is where readiness work begins.

START WITH DATA ACCESS AND PERMISSIONS

AI tools inherit the permissions of the accounts that use them. If a staff member connects Microsoft Copilot to their account, Copilot can access everything that account can access shared drives, email, documents, calendars.

In most organizations, file sharing has accumulated over years without intentional governance. Documents are overshared. Folders have permissions that were set during onboarding and never revisited. Sensitive files sit in locations accessible to far more people than necessary.

AI doesn't create this problem. It reveals it at scale.

> *AI can already see a lot more of your data than you think.*

WHAT TO DO BEFORE ENABLING AI

Ask your IT provider to conduct a permissions audit across your primary storage environments: OneDrive, SharePoint, Google Drive, Box, or whatever platforms you use. The goal is to identify overshared files and folders, reduce access to what's necessary, and document your data access structure before AI tools expand their reach.

This is not a one-time exercise. Permissions drift over time. Build a review cadence into your governance structure from the start.

BUILD A GOVERNANCE STRUCTURE

Sustainable AI readiness requires ownership. Someone or ideally a small cross-functional group needs to be responsible for how AI is used across the organization.

Effective governance structures include a working group with executive sponsorship, clear channels for staff to report concerns or request approvals, a feedback loop for sharing what's working, and alignment between IT, legal, and program leadership.

The governance structure doesn't need to be complex. It needs to be real. A policy without ownership is a document, not a program.

04 Establish Guardrails Before You Scale

Governance tells you who owns AI oversight. Guardrails tell everyone else what they're allowed to do. Both are necessary. Most organizations build one without the other.

THE AI ACCEPTABLE USE POLICY

An AI Acceptable Use Policy is the operational foundation of a responsible AI program. It doesn't need to be lengthy. It needs to be clear.

At minimum, an effective policy defines what data can and cannot be entered into AI tools, requires human review of AI-generated outputs before publication or use in decisions, establishes which AI systems are approved for use, and creates a process for requesting exceptions.

The policy should reference your existing Information Security Policy and align with applicable regulations relevant to your organization's work.

HUMAN OVERSIGHT IS NON-NEGOTIABLE

AI systems can assist and enhance human decision-making. They cannot replace it. Every policy should make this explicit: automated final decision systems are not permitted. Humans remain accountable for every output AI produces.

This isn't a limitation on AI's usefulness. It's the condition under which AI can be trusted.

MAINTAIN AN AI INVENTORY

Every AI system in use across the organization should be documented identifying each tool, the department using it, the data it accesses, and the staff responsible for oversight. The inventory should be reviewed annually and updated whenever new tools are introduced.

Most organizations are surprised by how many AI tools are already running when they conduct this exercise for the first time.

ANNUAL RISK ASSESSMENT

Each approved AI system should undergo an annual review covering security, privacy, legal, reputational, and competency risks. The AI landscape changes quickly. A tool that was appropriate at adoption may not be appropriate twelve months later.

> *Policy works better than prohibition. Culture matters more than technology.*

05 From AI Chat to AI Agents

Once governance and guardrails are in place, the question becomes how to scale AI use responsibly across the organization. The answer is a two-track approach.

TRACK ONE: BROAD PRODUCTIVITY ENABLEMENT

Enable general AI tools Microsoft Copilot Chat, Google Workspace Gemini, or similar platforms for broad staff use within the constraints of your Acceptable Use Policy. The goal here is familiarity.

Staff learn the fundamentals of prompt engineering. They discover where AI adds value in their daily work. They develop judgment about what AI does well and where it falls short. This track is about building organizational fluency.

TRACK TWO: ADVANCED PILOT

Select a smaller group to explore more advanced capabilities workflow automation, customized AI assistants, advanced integrations, mission-specific use cases. This group moves faster, tests more, and brings learnings back to the broader organization.

The pilot group is where you surface technical and compliance issues in a controlled environment rather than across the full organization.

IDENTIFY YOUR HIGH-IMPACT USE CASES

Before scaling, identify one to three use cases where AI can add meaningful value to your specific work. For nonprofits, common high-impact use cases include grant drafting and editing, internal knowledge search, policy summarization, and volunteer coordination.

For each use case, define clear success metrics, evaluate the risk level, and establish the oversight required. Use cases carry different risk profiles. A grant drafting assistant requires different oversight than an automated donor communication tool.

> *Select use cases based on value and risk not just what's technically possible.*

BUILD A PROMPT LIBRARY

As staff develop effective prompts for common tasks, capture them in a shared library. This accelerates onboarding, establishes best practices, and ensures the organization's AI use improves over time.

06 What Operationally Mature AI Adoption Looks Like

Operational maturity in AI adoption isn't about using the most advanced tools. It's about using AI in a way that's sustainable, accountable, and aligned with your organization's mission.

GOVERNANCE IS ONGOING, NOT ONE-TIME

Mature organizations treat AI governance as a continuous function. The steering committee meets regularly. The acceptable use policy is reviewed and updated as tools and regulations evolve. The AI inventory is maintained. Risk assessments happen on schedule.

Governance doesn't require significant overhead. It requires consistency.

STAFF ARE TRAINED AND SUPPORTED

AI literacy is built into onboarding and ongoing training. Staff understand what they're permitted to do, why the guardrails exist, and how to use AI effectively within them. There are clear channels to ask questions and request approvals.

Organizations that build AI literacy into their culture spend less time managing exceptions and incidents.

THE TECHNOLOGY ENVIRONMENT IS READY

Data permissions are clean. Integrations are documented and secured. Authentication is strong. The IT environment has been reviewed with AI in mind, not just general security hygiene.

This is the foundation everything else depends on. AI governance built on a poorly governed technology environment will not hold.

AI USE IS TIED TO MISSION OUTCOMES

The organizations that get the most value from AI connected it to specific mission outcomes from the beginning. They didn't adopt AI because it was available. They adopted it because they identified where it could help them do more of what matters.

Ready to take the next step?

Defensible helps organizations build the governance foundation that responsible AI adoption requires IT, cybersecurity, and policy.

defensible.tech/get-started