AI ADOPTION ROADMAP

A four-step guide to your AI journey.

A structured framework for responsible AI adoption — from foundational policy to measurable business outcomes.

PREPARED BY

KeyStone® solutions, llc

ITTaaS™ Managed Services · keystone.solutions

YOUR AI JOURNEY

Four steps. In order.

Most AI projects do not fail because the technology is bad. They fail because someone skipped step one. Work through the four steps below in sequence — each builds on the one before it.

01

STEP ONE · FOUNDATION

Foundational Principles

Build the four-part policy structure that governs how AI behaves in your organization — before you sign your first AI license.

02

STEP TWO · GOVERNANCE

AI Governance Framework

Decide who owns AI decisions, how new tools get approved, and how the policy stays current as the AI market shifts.

03

STEP THREE · AI IN USE

AI Tools Inventory

Build a living inventory of every AI tool in your organization, who uses it, and what state it is in.

04

STEP FOUR · OUTCOMES

Business Outcomes

Translate aspirations into measurable business outcomes with results, baselines, owners, and timelines.

01

STEP ONE · FOUNDATION

Foundational Principles.

Policy before tools. Generate your AI Usage Policy before you sign your first AI license.

STEP 1 OF 4 · FOUNDATION

Build the foundation.

Before deploying AI tools, build the policy foundation that governs how AI behaves in your organization. Bolting policy on after an incident is harder, more expensive, and misses the gap that caused the incident.

The four policies below are the structure of an AI Usage Policy. Decide on each one before you sign your first AI license. If you want a starting point, the policy builder at **policybuilder.keystone.solutions** generates a boilerplate draft based on your inputs — refine it to your business specifics and have qualified legal counsel review the final language before adopting it.

> ### Security-First by Design
>
> Decide your security posture before you select tools. Organizations that pick tools first end up writing policy to fit the tool, and the policy misses the right risks. Your applicable compliance framework (NIST CSF, HIPAA, FFIEC, or whatever governs your industry) sets the floor. AI policy extends it.

AI Security Posture

Your security posture sets the stance your organization takes toward AI risk. It defines the data sensitivity at play, the regulatory frameworks that apply, and the level of supervision required for AI output.

- The data categories your AI tools may and may not touch (PII, PHI, financial records, IP, customer data)
- The regulatory frameworks that govern your business and how they apply to AI use
- Whether AI output requires human review before it leaves the organization
- How AI use fits inside your existing access controls, logging, and incident response

AI Acceptable Use

Acceptable Use defines what employees may do with AI tools and what they may not. This section becomes a Permitted and Prohibited list. Decide what goes on each one.

- Use cases you want to permit (drafting, summarization, research, internal training material)
- Use cases you want to prohibit (inputting customer PII, generating client-facing communications, AI-driven decisions without human review)
- The approval path for use cases not on either list
- Whether contractors and vendors are bound by the same rules

Allowed AI Tools

Your Allowed Tools list is your approved AI inventory. Anything not on the list is not permitted. This section becomes a named list with scope of use for each tool.

- The AI tools you have evaluated and approved
- What each tool is permitted to do, and what it is not
- Who is authorized to use each tool
- The process for adding a new tool to the list

Compliance and Ethics

Compliance and Ethics assigns individual accountability and ties AI use to your regulatory and ethical obligations. This section becomes a numbered list of obligations every employee acknowledges.

- The reporting path when an employee observes a policy violation or unintended data exposure
- The retention requirements for AI-assisted work products
- The consequences for policy violations
- The standard for evaluating AI output before relying on it

02

STEP TWO · GOVERNANCE

AI Governance Framework.

Clear ownership and fast decisions. A named owner, a documented approval path, and a regular review cadence cover most organizations.

STEP 2 OF 4 · GOVERNANCE

Decide who owns it.

Governance defines who owns AI decisions, how new tools get approved, and how the organization keeps the policy current as the AI market shifts.

> ### Governance Without Bureaucracy
>
> The goal is clear ownership and fast decisions. Committees that meet every six weeks to debate a Copilot license fail this test. A named owner, a documented approval path, and a regular review cadence cover most organizations.

AI Owner

Name one person accountable for AI strategy, the approved tools list, policy updates, and vendor reviews. Without a named owner, the work happens by accident or not at all.

- A named individual with executive backing
- Authority to approve or reject new AI tools
- Ownership of the approved tools inventory and the policy itself
- A defined escalation path when a decision needs leadership input

Tool Approval Process

Document the path a new AI tool takes from request to approval. Without a written process, employees sign up for tools on their own credit cards, and your data goes with them.

- A request form or queue that captures business use case, data involved, and proposed user group
- A security and privacy review (vendor security certifications, data retention, training data use)

- A documented decision (approved, conditionally approved, rejected) with rationale
- A target turnaround time so requesters know what to expect

Training and Awareness

Employees cannot follow a policy they have not read. Set up onboarding training and an annual refresher tied to policy changes.

- Required AI awareness training before access to approved tools is granted
- Annual refresher covering policy updates and new tools added in the past year
- Role-specific training for high-risk users (finance, HR, legal, customer-facing roles)
- A documented acknowledgment that employees have read and understood the policy

Review Cadence

AI tools and vendors change faster than most policy cycles. Put reviews on a calendar at fixed intervals so they happen before something forces them.

- Quarterly review of the approved tools list and known shadow AI use
- Annual review of the full policy with IT, security, and leadership
- An incident review process that triggers an immediate policy revisit
- A vendor renewal review for AI tools coming up for contract renewal

03

STEP THREE · AI IN USE

AI Tools Inventory.

A living inventory of every AI tool in your organization. Build it once, then update it on a calendar.

STEP 3 OF 4 · AI IN USE

See what is already there.

A living list of every AI tool in use across your organization, who uses it, and what state it is in. Build it once, then review and update it on the cadence set in Step 2.

Three categories cover most AI in business use. If a tool spans more than one, list it under the category that matches its primary purpose.

Productivity & Content

AI tools embedded in productivity suites, drafting platforms, and content workflows. These have the broadest user base and the loosest data boundaries, which makes them the highest-priority inventory category.

TOOL / PLATFORM	USED BY	STATUS

Common examples: *Microsoft 365 Copilot, ChatGPT, Claude, Gemini, Notion AI, Grammarly. If employees use any of these, the tool belongs in your inventory whether or not your organization paid for it.*

Security & Operations

AI built into security tooling, monitoring platforms, and IT operations. These run inside your IT stack, configured by your MSP or internal IT team rather than end users.

TOOL / PLATFORM	USED BY	STATUS

Common examples: *Microsoft Copilot for Security, MDR platforms with AI correlation, AI-assisted SIEMs, RMM tools with AI summarization.*

Customer-Facing & Specialized

AI in customer-facing channels, or AI built for a specific business function (legal, finance, HR, sales). These tools carry the most legal and reputational risk because their output reaches outside the organization.

TOOL / PLATFORM	USED BY	STATUS

Common examples: *Website chatbots, AI-powered call center software, contract review tools, AI sales assistants, AI-powered HR screening platforms.*

Status definitions

Active — Approved, in production use, reviewed against the policy.

Managed — In use with conditions attached (limited user group, output review, scoped data access) to keep it inside policy.

Evaluating — Under review for approval. Not yet authorized for production use with sensitive data.

04

STEP FOUR · OUTCOMES

Business Outcomes.

A result the business cares about, a measurable target, a named owner, a timeline. Anything less is a topic.

STEP 4 OF 4 · OUTCOMES

Define what success looks like.

A business outcome names a result the business cares about, the metric that proves it happened, the person accountable, and the date you expect to see it. Without those four pieces, you have a topic.

Compare “AI for customer service” with “reduce average ticket resolution time from 18 hours to 6 hours by Q4, owned by the head of support.” The second example is a plan you can fund, staff, and review. The first is not.

What makes a real business outcome

A business outcome you can build a plan around has four things:

- A result the business cares about, stated in business terms (revenue, cost, time, risk, customer experience)
- A measurable target with a baseline (what is true today, what will be true if the work succeeds)
- A named owner with the authority to make the call when something needs to change
- A timeline tied to a real decision point (budget cycle, quarterly review, renewal date)

Without all four, the work is not ready to fund.

Common outcome categories

The list below shows common categories. Yours may include others. Most organizations find their real outcomes by starting from a top-three operational pain point and working backward into AI capabilities.

01 Operational Cost Reduction

AI replacing manual effort in repetitive admin work: scheduling, drafting routine communications, meeting notes, document generation, data entry. Measurable in hours saved, headcount avoided, or process cost per transaction.

02 Customer Experience Improvement

AI-assisted customer service: knowledge-base-connected assistants for first-contact questions, ticket routing, response drafting for human review. Measurable in resolution time, deflection rate, customer satisfaction, ticket volume.

03 Business Intelligence and Reporting

AI that pulls data from multiple sources, identifies trends, and produces narrative reports. Measurable in time-to-insight, report cycle time, and the number of decisions made on AI-prepared analysis.

04 Security and Risk Posture

AI-assisted threat detection, anomaly identification, and incident triage. Measurable in detection time, false-positive rate, mean time to contain, and the value of incidents prevented.

05 Onboarding and Workflow Acceleration

AI-guided onboarding for new employees or new customers. Measurable in time-to-productivity, error rate on provisioning, support tickets opened in the first 30 days.

06 Document and Contract Intelligence

AI that summarizes contracts, flags non-standard clauses, extracts key dates and obligations, and accelerates document review. Measurable in review cycle time, missed renewal dates, legal hours per contract.

Your outcomes

The six categories above cover the most common ground. Your actual outcomes will be specific to your business, your industry, and the pain points you have today. Pick the two or three that matter most this year. Write them down as real outcomes (result, baseline, target, owner, timeline) before you spend a dollar on an AI tool to deliver them.

USING THIS ROADMAP

Work through the four steps in order.

Policy first. Governance second. Inventory third. Outcomes last. Reversing the order is the most common pattern in AI projects that fail.

This document is a guide. The next step is a written AI Usage Policy your team will actually follow. Start one from scratch, or use the policy builder at **policybuilder.keystone.solutions** to generate a boilerplate draft. Either way, refine it to your business specifics and have qualified legal counsel review the final language before adopting it.

TRY THE POLICY BUILDER →