

A PRACTICAL GUIDE

AI Tooling

A strategy for the tools, not a bet on one.

01 Your AI Tool Portfolio

Your company is already using AI tools. The question is whether anyone decided which ones.

Marketing is using ChatGPT. Engineering tried GitHub Copilot. Someone in operations discovered Claude and has been quietly using it for three months. The CEO saw a demo at a conference and asked the team to "look into it." Nobody coordinated. Nobody compared. Nobody asked about the licenses.

This is not unusual. Most organizations adopted AI tools the same way they adopted smartphones for work: individually and informally. Faster than any policy could keep up with. When smartphones started accessing company email and documents before MDM policies existed, organizations learned that adoption without governance creates risk. AI tools carry the same lesson at higher stakes. They process your business data, your strategy documents, your financial information. The question is not whether your team is using them. It is what happens to the data when they do.

The organizations getting real value from AI are not the ones that picked the best tool. They are the ones that stopped treating tool selection as a personal preference and started treating it as a business decision. Not a one-time vendor pick, but an ongoing strategy: what tools, for what work, under what rules, with a plan for when things change.

This guide walks through the decisions that matter, the risks most organizations are not tracking, and how to build a tool strategy that does not lock you in.

Key Terms

These terms come up throughout the guide. This page is a reference you can return to.

Prompt	Your input to the model. Everything you type into the conversation is a prompt. The quality of the prompt directly determines the quality of the output.
Model	The AI intelligence underneath any tool. GPT, Claude, and Gemini are models. A model takes your input, reasons over it, and generates a response. You do not interact with a model directly. You interact through a tool.
Tool / Harness	The interface you work in. ChatGPT, Claude Desktop, Microsoft Copilot, and Cursor are all tools. Each runs on one or more models. The tool shapes the experience. The model shapes the quality.
Token	The basic unit of text that AI models work with. Not quite a word, not quite a character: somewhere in between. "RoleModel" might be two tokens ("Role" and "Model"). When you hear about model costs or context limits, this is the unit being measured.
Context window	The amount of text the model can see at once during a conversation. Think of it as the model's working memory. Everything inside the window is visible and can inform the response. Everything outside it is invisible.
Hallucination	When a model generates something plausible but incorrect. It is not lying. It is predicting what a good answer looks like. Most common on specific facts, numbers, and anything requiring precise recall. The practical response: verify before it leaves your desk.

02 The Tools Today

The AI tools your organization encounters fall into three categories. The categories matter more than the product names, because the products change every few months.

Chat

ChatGPT. Claude. Gemini. The front door most people walked through first. You open a browser or desktop app, type a question, and get an answer. Good for exploration, drafting, Q&A, and working through ideas. The defining characteristic: you bring your work to the AI. It works with what you give it in that conversation and nothing more.

Embedded AI

Copilot in Microsoft 365. Gemini in Google Workspace. AI features inside Slack, Notion, your CRM, and industry-specific tools. The AI comes to you, already inside the tools your organization pays for. The appeal is convenience. The limitation is isolation: Copilot in Word does not know what Copilot in Excel is doing, and these features often get enabled without anyone reviewing what data the AI can access.

Agent Tools

Claude Code, GitHub Copilot (the coding agent, distinct from the chat feature), Cursor, and a growing set of tools that connect to your systems and take action across them. They read files, run commands, pull data from connected systems, and execute multi-step workflows. Increasingly relevant beyond software development for anyone doing work that spans multiple systems.

03 The Data Question

Before you evaluate features, compare pricing, or choose a category, there is a question most organizations skip: **where does your data go?**

Every AI tool your team uses processes your inputs. What happens to those inputs after you hit send depends on the tool, the plan tier, and the terms your organization agreed to (or never read). This section is the one worth forwarding to your IT lead or compliance team.

Free, Personal, and Business Plans

The plan tier your team is on determines what happens to your data. The differences are not cosmetic.

- **Free tiers** train on your inputs by default. Retention policies are opaque or nonexistent. No admin visibility, no audit trail, no way for your organization to manage the account. No IP protections for what the AI produces.
- **Personal paid plans** (ChatGPT Plus, Claude Pro) are better but still consumer-grade. The provider typically does not train on your inputs. But there is no admin control, no centralized billing, no organizational visibility. No IP indemnification. The account belongs to the individual. When that person leaves, the account goes with them.
- **Team and business plans** (ChatGPT Team/Enterprise, Claude Team/Business) are where organizational control begins. No training on your data. Configurable retention. Audit logs. Admin controls. Centralized billing. IP indemnification for AI-generated content. The organization owns the accounts, not the individuals.

The gap between personal paid and team plans is the one most organizations miss. A team where everyone has their own ChatGPT Plus subscription feels like it has solved the AI tool question. It has not. There is no admin visibility, no centralized governance, no IP protection, and no way to revoke access when someone leaves.

Shadow AI

Research shows over 80% of workers use AI tools their employer has not approved.¹ More than half enter sensitive data: financial information, strategy documents, and details specific to the work they do for their organization.

This is not a discipline problem. It is a supply problem. People use personal accounts on company work because the company has not provided an approved tool on a business plan. The fix is not a policy memo telling people to stop. It is providing the approved tool before they find their own.

¹ Microsoft and LinkedIn, *2024 Work Trend Index Annual Report*. Salesforce's 2024 survey of 14,000 workers found similar rates.

An AI Policy Is Not Optional

"Use AI responsibly" tells your team nothing. It does not answer the questions people actually have. A usable AI policy answers concrete questions:

- Which tools are approved for business use?
- What data can go into them, and what cannot?
- What plan tier is required, and who pays for it?
- Can employees use personal accounts for work tasks?
- Who decides when a new tool gets added or removed?
- What happens to accounts and data when someone leaves?
- Have you reviewed data access for embedded AI features you already pay for?
- Is there a named owner of AI policy, not just "IT" or "leadership"?
- Are you tracking AI-specific regulations in your operating jurisdictions?

This does not need to be a 40-page document. It needs to be specific enough that someone can read it and know whether they can paste a contract into ChatGPT.

04 What Matters in Choosing

Once you know where your data goes, two questions drive the actual tool decision.

Match the Capability to the Work

Most AI subscriptions today give you access to a chat interface, a desktop app, and for some, coding and agent tools. The question is not which product to buy. It is which capability to reach for.

Quick drafts, Q&A, and thinking through ideas: the chat interface handles this well. Document editing, email summaries, and in-app convenience: embedded AI is built for this. Multi-step workflows, connected data, repeatable processes: this is where agent tools earn their place. Match the capability to the complexity of the work.

Lock-in Is the Real Risk

The model itself is compute you are renting. The provider is not an ally in this. They are a utility. Pricing will change. Open source models will become viable for more workflows. Today's best option will not be tomorrow's.

The question to ask about every tool you adopt: **if this tool doubles its price or a better option appears, how much work is it to move?**

The lock-in principle

Put your skills in files you own. Store your workflows, prompts, and project context in documents and repositories your organization controls. Some tools offer project-level containers (like Claude Projects) that keep context organized within the tool. Those are useful, but they are not yours if you leave. The portable version lives in your files. When the model is rented compute and the workflow is yours, the tools can change without your strategy breaking.

05 Model Choice

The Team Analogy

Most AI providers offer a tier system. The tiers map to a pattern you already know: a team of people at different levels.

SENIOR	The Senior on a Hard Problem Thorough reasoning. Slower. Worth the time on high-stakes or genuinely complex work. Your escalation, not your default.
DEFAULT	Your Daily Default Handles the large majority of professional work. Drafting, analysis, research synthesis, structured documents. Start here.
ANALYST	The Analyst on a Quick Task Fast and lightweight. Summaries, reformatting, simple lookups. You would not assign a senior consultant to clean up a spreadsheet.

This pattern exists across every provider: ChatGPT, Claude, Gemini, and others.

The practical heuristic: Start with the default. Escalate when the output is not meeting the bar. Every major provider follows the same pattern.

	Analyst	Daily Default	Senior
ChatGPT	Instant	Thinking	Pro
Claude	Haiku	Sonnet	Opus
Gemini	Flash	Pro	Ultra

06 Working Across Tools

No single AI tool does everything well. The organizations getting the most from AI are not picking one winner. They are combining tools based on what each does well right now.

Research and synthesis

One tool for broad, autonomous research. Another for disciplined synthesis. The research tool produces breadth; the synthesis tool produces discipline. Kick off a research run, then hand the results to a different tool with clear instructions on what to produce.

Visual generation

Some tools generate images: illustrations, editorial visuals, stylized graphics. Others produce structured diagrams, tables, and layouts you can edit and refine. Neither is universally better. Run the same concept through both and pick the one that fits the context.

Adversarial review

Paste the output from one tool into a different tool and ask for critical pushback. A different model, with no investment in the prior answer, will see the gaps differently. Ask: What is the weakest claim? What is missing? Where does the language do work the evidence does not? The output will be uncomfortable. That is the point.

Parallel drafting

Give the same brief to two different tools and compare the outputs. One may structure the argument better. One may catch a nuance the other missed. One may produce a tone that fits the audience. You are not picking a winner. You are pulling the best elements from each into a final version that neither tool would have produced alone.

07 The Handoff

When you move between tools, context does not follow automatically. The skill that makes multi-tool work effective is a structured handoff. Paste this prompt into your current tool before you switch. Copy the output into your next session.

PROMPT

Context Checkpoint

Before this conversation ends, produce a Context Checkpoint. Use this structure:

Decisions made: Three to five sentences on what we have decided and why. Focus on the decisions that matter, not every step we took.

Active constraints: What is fixed, what cannot change, what trade-offs we have already locked in.

In flight: What we are working on right now, and what state it is in.

Next steps: A clear, numbered list of what comes next. Someone reading only this should know exactly where to pick up.

Critical references: Three to five things the next session needs: a requirement, a key example, a data point, a document excerpt, or a formatting constraint.

Keep the whole thing under 400 words. Lean toward clarity over completeness.

The 400-word limit forces distillation. Without it, most people either paste too much context (which the next tool ignores) or provide too little (which produces generic output). This works across any combination: ChatGPT to Claude, Claude to ChatGPT, or the same tool in a new session. Once the habit is established, save this prompt as a skill in your preferred tool so you can invoke it by name.

08 Building a Tool Strategy

Moving from accidental adoption to a deliberate strategy does not require a six-month planning cycle. It requires four concrete decisions.

Audit What You Have

Before choosing new tools, know what your organization is already using. This audit is often surprising. Who is on free tiers? Who is paying for personal plans out of pocket? What data is going into each tool? Are there tools with overlapping purposes that nobody coordinated on?

Decide on Tiers

Not every role needs the same tools.

Baseline. An approved chat tool on a business plan, available to everyone. Covers exploration, drafting, Q&A. This is the minimum: an approved option that keeps people off personal accounts and free tiers.

Power users. Agent tools for people doing complex, multi-step work. Proposal writers, analysts, project managers, engineers. These roles benefit from tools that connect to systems and execute workflows.

Team systems. Shared skills and connected tools for workflows that span people and departments. This is where the investment pays off at organizational scale.

More in this series

Guide 2: AI Skills covers how to encode your team's expertise into reusable workflows that any AI tool can follow.

Guide 4: The AI Workspace covers how to build a connected system around agent tools that your organization controls.

Set the Rules

Write an AI policy that answers the concrete questions your team actually has. Approved tools. Data boundaries. Licensing requirements. The process for evaluating and adding new tools. What happens to accounts when someone leaves.

Keep it short enough that people will actually read it.

Plan for Change

Lock-in is the silent cost of every tool decision. The model is rented compute. Build your workflows in portable formats. Encode your expertise in skills and processes that live in files you own, not in features that live inside a vendor's app.

When the pricing shifts, when a better model ships, when open source becomes viable for a workflow you care about, you want to adapt. Not rebuild.

NEXT STEP

Free Workflow Assessment

We start with the workflow, not the technology. We look at how the work actually gets done today, where the friction is, and come back with two things: what you could do on your own, and options for how RoleModel could partner with you to build it.

Our first recommendation may be a Discover Phase, a short engagement that maps your specific workflows, team, and use cases to the right path for AI enablement. Process first, software second.

[Schedule a Free Workflow Assessment](#)