

# Practical AI for Organizations.

A starter briefing.

RYEXDEV · MAY 2026

# What this briefing is.

A small, opinionated set of slides for leaders, operators, and teams trying to put AI to work.

It covers four things:

- The AI landscape today, in plain language.
- A five-rung ladder for thinking about how AI gets used inside an organization.
- What's safe to do with your data, and how to decide.
- How to roll AI out to a team without it turning into chaos.

# The five-rung ladder.

A way to think about where you are.

---

1. Just open the chat. Free tools, talked to like a smart colleague.
2. Bring your data in. Files, photos, documents in the loop.
3. AI inside the tools you already pay for. Office and Workspace.
4. Agentic. AI doing twenty things on its own before reporting back.
5. Custom systems and orchestration. Specialized agents in parallel.

Most organizations get most of their value from rungs 1 to 3. Slow, careful experiments at rung 4. Rung 5 is where leverage is, but the cost to build is real.

# The major AI carriers in 2026.

Four cover most general needs. Four more cover specific lanes.

## Anthropic – Claude

Careful writing, coding, deep document analysis, agentic work. 1M token context standard.

No image, video, or voice generation.

## OpenAI – ChatGPT family

Broad general use. Voice mode, vision, image generation, transcription via Whisper.

Sora video being sunset April–September 2026.

## Google – Gemini family

Widest coverage. Chat, video (Veo), image (Imagen), uploaded-source synthesis (NotebookLM), deep Workspace integration.

## Microsoft – Copilot

AI inside Office. Agent Mode now acts inside Word and Excel. Strong governance via Purview.

# The specialists.

For when you need one thing done well.

## Perplexity

AI search with cited sources. Audit-friendly research answers.

## ElevenLabs

Voice generation, cloning, multilingual narration. The pick for any project that needs human-sounding voice.

## Midjourney

High-aesthetic image generation. Marketing visuals, concept art, brand imagery.

## NotebookLM (Google)

Synthesize uploaded sources into briefs, audio overviews, and infographics. Bookmark-worthy for research-heavy work.

# How the buckets get covered.

BUCKET	REACH FOR
Chat, writing	Claude, ChatGPT, Gemini
Coding	Claude, ChatGPT
Document and data analysis	Claude, ChatGPT, NotebookLM
Agentic work	Claude (Code), Microsoft (Agent Mode)
Web research with citations	Perplexity, Gemini, ChatGPT
Image generation	Imagen, ChatGPT image, Midjourney
Voice and audio	ElevenLabs, ChatGPT voice mode, Whisper
In-product workflows	Microsoft 365, Google Workspace

Every bucket has at least one strong option. No one carrier covers it all.

# What's safe to do.

Three rules that hold up under audit.

---

The billboard rule. Treat any input to a public AI as if it could appear on a billboard tomorrow. Anything proprietary needs an enterprise plan with a written zero-retention agreement, or it stays out.

The permissions rule. Apps on your devices only access what you grant them. If a tool wants full-disk or full-camera-roll access to "help," that's the access it has. Audit permissions yearly. This is true for any app, AI or otherwise.

The tier rule. Paid consumer tiers (ChatGPT Plus, Claude Pro) usually don't train on your inputs by default, but defaults change. Enterprise plans (Team, Enterprise, API) put that promise in writing. If your work touches anything proprietary, choose the tier that puts the promise in

# How to trust the output when it matters.

For any output where mistakes have a cost.

---

1. Source the answer. Ask the AI to cite. If you can't trace the claim back to something, don't treat it as ground truth.
2. Constrain the data. Don't ask the AI to remember facts. Hand it the catalog, the spec, the standard, and ask it to find the answer in what you provided.
3. Build verification into the workflow. At scale, the check has to be in the system, not in someone's head. Schema validation, automated spot-checks, human review on sampled output before publish.

Every time AI burns an organization, the same thing has happened: the output looked right, was not, and nobody had built the check.



# Rolling AI out to a team.

A version that holds up.

---

- Pick one default tool your organization is licensed for, with proper security and retention. People use it unless there's a specific reason not to.
- Allow named alternatives for named tasks. Perplexity for research, GitHub Copilot for engineers, a specialized model for a specialized workflow. With a written rationale.
- Write the policy down. One page. What's allowed, what's not, what to never paste anywhere, how to report a near-miss. Mandate it.
- Train hard at the bottom of the ladder. Most people in your organization will never go past rung 2. Make rungs 1 and 2 dead-comfortable for everyone.

# When DIY stops being the right trade.

The places where doing it yourself stops being the right call.

- Validation at volume. Hand-checking AI output at fifty records works. At fifty thousand, you need engineering.
- Schema-bound work. ACES/PIES, FHIR, regulated disclosure formats. "Looks right" is not the bar.
- Integration with legacy systems. ODBC, SOAP, EDI, flat files. Wiring AI into those reliably is its own skill.
- Governance and security posture. Data residency, retention, audit logs, vendor risk. Real once AI is part of operations.

If three or more describe your situation, you're in territory where the right path depends a lot on your specific systems, data, and team.

## Where to go next.

This briefing lives alongside two other artifacts in the same kit.

---

- AI-Use-Policy-Outline.docx – a starting outline for HR/Legal to adapt into an actual policy for your organization.
- Starter-Prompts.md (and .pdf) – ten prompts organized by role. Drop into any AI tool to get started.

The longer-form thinking lives at [ryexdev.com/ai](https://ryexdev.com/ai).

RYEXDEV · MAY 2026

# Questions? Reach out.

[ryexdev@gmail.com](mailto:ryexdev@gmail.com)

[ryexdev.com](http://ryexdev.com)