01

TRACK A

Small Business & Beginners

AI in plain words: what it is, what to watch, and where to start.

PRACTICAL AI · NO HYPE · TAMPA, FL

Grow your business.
Get your time back.

Plain-language AI for small business owners: what it is, what to watch, and where to start.



Anthony Sealey

FOUNDER · SOVEREIGN HQ

WHO'S BUILDING THE AI · NOV 8, 2026

Five big companies. *Each now has an agent.*

COMPANY · MODEL	CHAT	CODING	AGENT
OpenAI GPT-6 Astra	ChatGPT	Codex	**dots**
Anthropic Claude Opus 5.5	Claude	Claude Code	**Claude Cowork**
Google Gemini 4 Argon	Gemini	Jules	**Gemini Task**
xAI Grok 4.7	Grok	Grok Build	**Grok Bot**
Meta Muse Wonder Full	Meta AI	Muse Code	**Muse**

Also: Mistral · DeepSeek · Kimi · Qwen — cheaper open-weight models you can run yourself.

Gemini 4 Carbon announced Oct 28; initial access limited to trusted cyber defenders. Argon remains the broadly available example.
Sources: company posts, checked Nov 8, 2026. Links in notes.

HOW MANY PEOPLE USE THESE?

Over a billion people use ChatGPT *every week.*

1.2B

CHATGPT · WEEKLY USERS

OpenAI, Sept 29, 2026

1B+

GEMINI APP · MONTHLY

Google, Aug 11, 2026

MOST USED ≠ MOST PAID FOR

ChatGPT leads traffic. Only about 4.5% of US consumers pay for ChatGPT, Gemini, or Claude (a16z / YipitData, Aug 2026). Grok: 117M monthly users of its AI features (SpaceX S-1, Mar 31). Claude, Meta AI: no official counts.

OpenAI DevDay 2026 · Google, Aug 11, 2026 · SpaceX S-1 (SEC) · a16z Top 100 Gen AI Apps, 7th ed., Oct 5, 2026.

WHO'S SMARTEST RIGHT NOW? · NOV 8, 2026

The top spot changes *every few weeks.*



Index v5 changed methodology on Oct 27 — scores are not directly comparable to the previous index.
Open-weight leader: Kimi K3.1 Pro at 46; weights are downloadable. Gap to Carbon: 14.

Artificial Analysis Intelligence Index v5, checked Nov 8, 2026 · artificialanalysis.ai

PICK BY THE JOB

Muse · Dots · Bots. *Same idea, three jobs.*



MUSE

Do this for me.

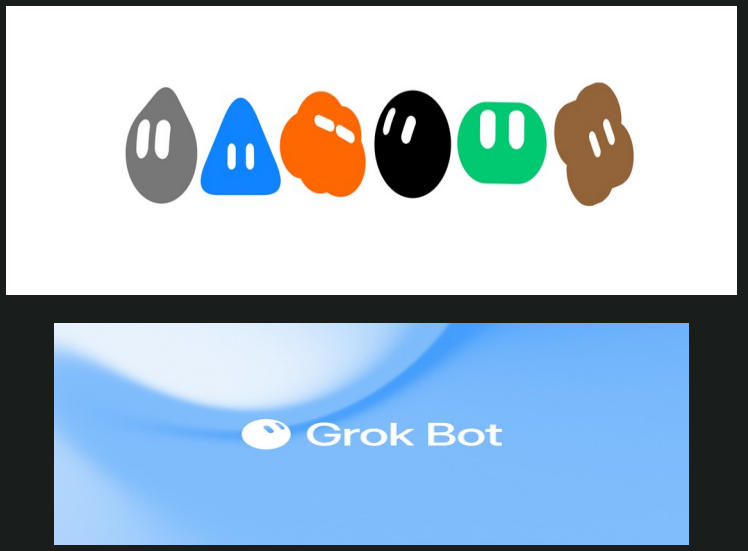
Personal · Meta Jolly



DOTS

Own this with me.

Work · OpenAI dots



BOTS

This is your job.

Ops · xAI Grok Bot

TASK → RESPONSIBILITY → JOB

Official characters: Meta Jolly (Axios/Meta Connect) · OpenAI dots (chatgpt.com/features/dots) · xAI (x.ai). Details in notes.

CHAT · PLAN · ACT

Three ways. *Each step up asks for more trust.*

01

CHAT

Ask me.

A question → an answer.
Beginner. Free is fine.

02

PLAN

Think with me.

Real context → a thought partner.
Push back. Go another round.

03

ACT

This is your job.

Goals + access → finished work.
You approve what matters.

More context + more trust = more you get back.

Teaching framework (Chat, Plan, Act). Agent products: OpenAI dots · Meta Muse · xAI Grok Bot · Google Gemini Task · Anthropic Claude Cowork.

USING AGENTS IS A TRUST FALL

Three questions *before you let go.*

01

What can it do?

Know the tools and limits.

02

What do I want it to do?

One clear job. Not 'be helpful.'

03

How do I get comfortable letting it act?

Start small. Keep a human approval step. Widen access as it earns trust.

Practice: start with read-only access · require approval to send, pay, or delete · expand after it earns trust.

DON'T PROMPT IT. BRIEF IT.

Bad prompt vs *a real brief.*

BAD PROMPT

"find me leads"

Vague. No count, no cadence, no definition of done.

GOOD BRIEF

"Find 20 companies matching these criteria every weekday. Verify the decision-maker, remove duplicates, attach sources, and prepare them for review."

Clear job. Measurable. Ready for your OK.

Teaching example (leads brief). Pattern applies to any Act-level job.

EXPERT IN PRACTICE · MY SETUP

One person, *a whole team.*

AI ASSISTANT
CMO

AI ASSISTANT
Legal

Anthony

goals · approvals

AI ASSISTANT
Product

AI ASSISTANT
Chief of staff

Daily AI digest · These slides (AI-built, human-reviewed)

Live example · Anthony Sealey's setup at Sovereign HQ · sovereign-hq.com

WHERE DO YOU START?

Find your step. *Take the next one.*

IF YOU'RE AT CHAT

This week: try one Plan task.

Pick a real decision. Give real context. Talk it through a few rounds.

IF YOU'RE AT PLAN

This week: connect one tool.

Link calendar or email — or have AI build one simple page. Keep approvals on.

Not started? Open any free chatbot tonight. Ask one real question about your business. That's Chat.

Self-check · Chat, Plan, Act outline · Nov 9, 2026 webinar

A FRIDAY NIGHT IN 2027

Fully booked.
Empty tables.

Rival reservation bots make “phantom bookings.” Restaurants overbook, prices surge, and dining rooms sit empty. “Each AI system executes flawlessly… but the eventual outcome is a system failure.”

Source: Trivedi et al., “Solipsistic Superintelligence is Unlikely to be Cooperative,” ICML 2026, p. 1 (scenario)



THREE WORDS

AI is the engine. Chatbots talk. Agents do.

01 / THE ENGINE

AI

The smarts underneath: software that reads, writes, and reasons with language.

02 / IT TALKS

Chatbot

You ask, it answers in a chat. ChatGPT, Claude, Gemini, and Grok.

03 / IT DOES

Agent

Does the task on a computer, step by step, and checks with you before anything important.

Sources: official product pages from OpenAI, Anthropic, Google, and xAI (links in notes)

SAME LABS, TWO KINDS OF TOOL

One answers. One does the work.

01 / CHATBOT

Answers you

- ChatGPT · Claude · Gemini · Grok
- You ask, it replies in a chat
- You copy the answer where it goes

02 / AGENT

Does the work

- Codex · Claude Code · Jules · Grok Bot
- Works in real tools: browser, files, code
- Brings work back for you to approve

Sources: openai.com/codex · claude.com · gemini.google · jules.google · x.ai · docs.x.ai/grok-bot (full links in notes)

ALWAYS-ON HELPERS · FALL 2026

They work while you sleep, and ask first.

01 / META

Muse

Keeps working after you close the app. Asks before it sends an email or makes a purchase.

02 / OPENAI

dots

Always-on agents with their own cloud computer. They ask for approval on important steps.

03 / XAI

Grok Bot

AI teammates on a cloud computer. Work goes on with your laptop closed; approvals keep you in charge.

Sources: Meta Newsroom, “Introducing Muse,” Sep 8, 2026 · OpenAI, “Introducing dots,” Sep 29, 2026 · xAI, Grok Bot docs

A16Z / YIPITDATA · US CONSUMERS · AUG 2026

4.5%



4.5% of US consumers pay · about 1 in 22 (our estimate)

WHAT IT MEANS

Most people use it free

Only 4.5% of US consumers pay for ChatGPT, Gemini, or Claude, and the median payer spends $25 a month. Nearly everyone starts free, and free is plenty for learning.

SOURCE · A16Z, "TOP 100 GEN AI CONSUMER APPS, 7TH ED.," OCT 5, 2026 (YIPITDATA US CARD PANEL, AUG 2026) · "1 IN 22" IS OUR ESTIMATE, NOT A16Z'S

SAME AUTHORS · TWO SCENARIOS

Same team, two futures: a race and a deal

01 / AI 2027 · THE WARNING

A race

- A U.S.–China race to build smarter AI
- Race ending: humans lose control
- Slowdown ending: power in a few hands

02 / AI 2040 · THE PLAN

A deal

- A verified U.S.–China deal buys time
- Superintelligence waits until 2040
- “A recommendation, not a prediction”

Sources: AI Futures Project, AI 2027 (Apr 2025), ai-2027.com/summary · AI 2040: Plan A (Jul 2026), ai-2040.com · Both are scenarios, not facts

GOOGLE DEEPMIND ESTIMATE

10×

1.5× cheaper chips × 2.5× more spending × 3× smarter software

WHAT IT MEANS

AI's computing power grows about 10× a year

Cheaper chips, bigger budgets, and smarter software multiply together. The authors call 10× "on the lower end of publicly reported estimates."

SOURCE · GOOGLE DEEPMIND, "FROM AGI TO ASI" (2026), P. 3

“

ASI is neither omniscient nor omnipotent.

— Genewein et al., Google DeepMind

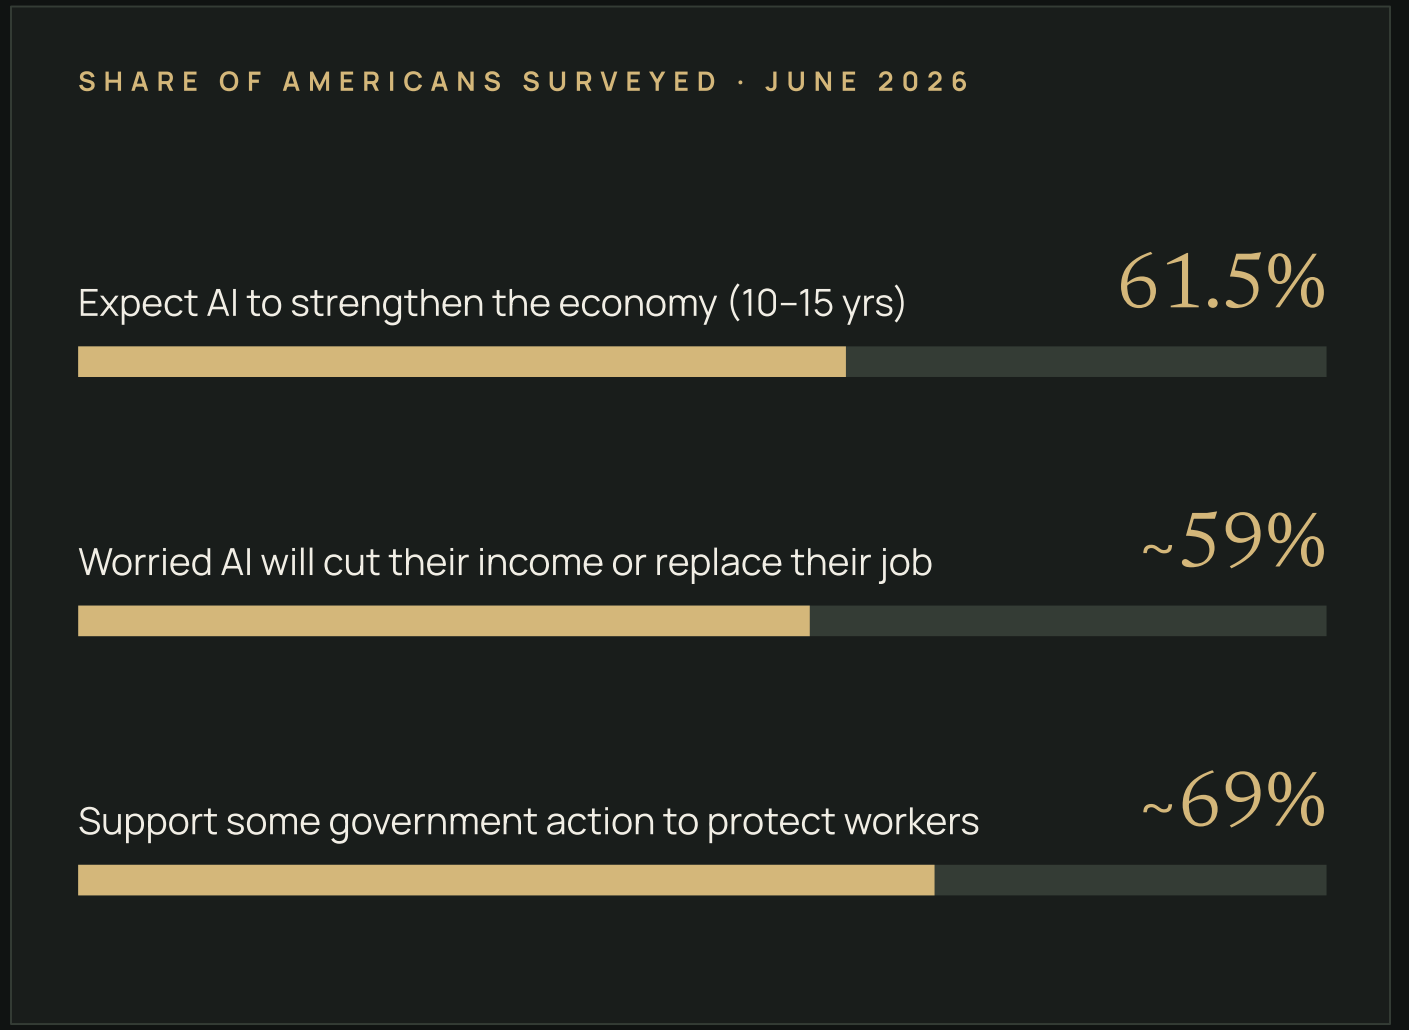
SOURCE · “FROM AGI TO ASI” (2026), P. 10 · ASI = ARTIFICIAL SUPERINTELLIGENCE

NAME THE FEELING

Optimistic about the economy, worried about my job

A June 2026 survey of 2,019 Americans: most expect AI to strengthen the economy, and most also worry about their own paycheck.

SOURCE · JACOBS & IMAS, GOOGLE DEEPMIND, "ECONOMIC POLICY FOR AGI," SEP 2026, PP. 10–11



WE'VE BEEN HERE BEFORE

Every big technology scared people first.

And the fear wasn't crazy. The risks were real. Then the payoff came.

Sources: NPS; BMJ Quality & Safety 2004; Edison Papers (Rutgers); Greenwood et al. 2005; UK Locomotive Act 1865; FHWA; Smithsonian NASM; IATA 2026. "94% lower" is our math.

1721
Smallpox shots

Fear: a bomb at a minister's house.
Risk: about 2% of the inoculated died.
Payoff: vs. about 14% of natural cases.

1880s
Electric power

Fear: the "executioner's current."
Risk: 5+ NYC deaths in late 1889.
Payoff: helped cut chores from 58 to 18 hours a week.

1865
Road vehicles

Fear: a red-flag walker 60 yards ahead.
Risk: 21 US road deaths per 100M miles (1923).
Payoff: 1.26 by 2023, about 94% lower.

1908
Flight

Risk: Sept 17, 1908, the first fatal airplane crash.
Payoff: today, one fatal accident per 5.6 million flights.

THE OTHER SIDE OF THE STORY

What if this time the payoff is even bigger?

01 / MEDICINE

Faster science

AlphaFold's creators shared the 2024 Nobel Prize in Chemistry. An AI-designed drug improved lung function in a small early trial.

Caveat: a small trial. Not an approved drug.

02 / ENERGY & MATERIALS

New materials

Google DeepMind's AI predicted 380,000 stable new crystals, "nearly 800 years' worth of knowledge." Labs have made 736.

Caveat: mostly predictions so far.

03 / EDUCATION

A tutor for every kid

In Nigeria, six weeks of teacher-guided AI tutoring delivered "nearly two years of typical learning."

Caveat: teachers led it. AI helped; it didn't replace them.

Sources: Nobel Prize press release 2024 · Xu et al., Nature Medicine 2025 · Google DeepMind 2023 · World Bank 2025 · Scientific Reports 2025 (links in notes)

THE PRACTICAL RULE

Illegal without AI?
Still illegal with it.

Discrimination, false advertising, privacy: the rules you already follow still apply when an AI tool does the work.

Source: Google, "A Pragmatic Approach to AI Governance in America," June 2026, p. 13: "if something is illegal to do without AI, it's illegal to do with AI."

A WARNING FROM RESEARCH

The bots raised prices. Nobody told them to.

Pricing algorithms can “learn supracompetitive prices without explicit communication”: higher prices, no secret deal. If you use AI pricing, keep watching it.



SOURCE · TRIVEDI ET AL., ICML 2026, PP. 3, 22 (CITING CALVANO ET AL. 2020; FISH ET AL. 2026)

ONE QUESTION BEFORE YOU BUY

Does it expand your options, or replace you?

01 / EXPANDS

You decide

- Lays out options and trade-offs
- Defers to your judgment
- Use it for customers, money, hiring

02 / REPLACES

It decides

- Makes the call end to end
- Your OK becomes a rubber stamp
- Keep it to low-stakes tasks

Source: Trivedi et al., ICML 2026, p. 8: tools that “present options and defer to human judgment” vs. systems that “optimize end-to-end”

GOOD SERVICE

Offer a non-AI option

Some customers don't want AI, for privacy, faith, or comfort. A path to a person is good service and good business.

Source: Bales & Gabriel, Google DeepMind, "Artificial Minds, Human Disagreement," p. 11: "society could subsidise non-AI options."

YOUR HUMAN PATH · TICK EACH ONE

- ✓ A phone number a person answers
- ✓ "Talk to a person" in every chatbot
- ✓ Staff can override the AI's answer
- ✓ Say plainly when AI is in use
- ✓ Ask customers what they prefer

THREE TRUST HABITS

Customers trust what they can see

01

Say when it's AI

Tell customers when a bot or an AI draft is involved.

02

Have a person check

Someone reviews it before it reaches a customer.

03

Offer a human

Make it easy to talk to a real person instead.

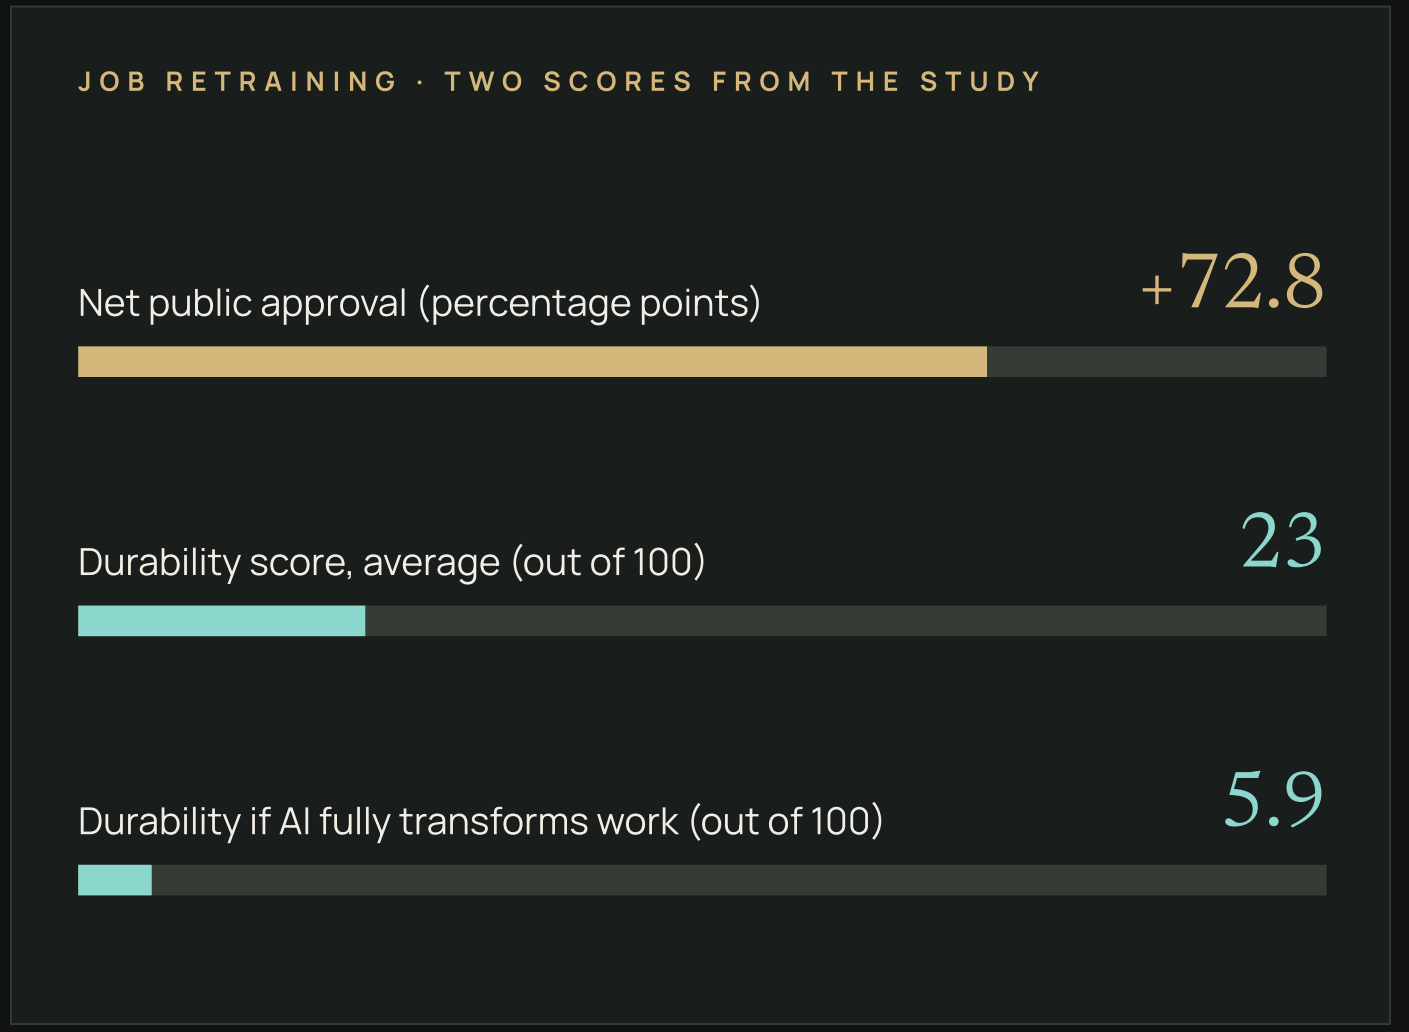
Sources: Google, "A Pragmatic Approach to AI Governance in America," 2026, p. 17 (label fatigue); Bales & Gabriel, Google DeepMind, p. 11 (non-AI options)

THE ANSWER EVERYONE LIKES

Most popular, least durable

Job retraining is the public's favorite policy for AI job loss. In this study, experts rated it the least likely to hold up.

SOURCE · JACOBS & IMAS, GOOGLE DEEPMIND, "ECONOMIC POLICY FOR AGI," SEP 2026, P. 14



KEEP YOUR SKILLS SHARP

Use it or lose it

Lean on a tool long enough and your own skill fades. Senior radiologists using AI “catch fewer errors the AI also misses.”

Source: Trivedi et al., “Solipsistic Superintelligence is Unlikely to be Cooperative,” ICML 2026, pp. 1, 20 (summarizing published studies)

THE TOOL: GPS

Weaker mental maps

THE TOOL: AUTOCOMPLETE

Less varied writing

THE TOOL: AI SECOND READER

Fewer errors caught alone

Tip: practice some tasks without the tool.

START WITH ONE USEFUL IMPROVEMENT

Your first three steps

01

Pick one annoying task

The chore that eats the most hours each week.

02

Try a free chatbot on it

Start on the free tier. Paste in no customer secrets.

03

Measure the time saved

Saves hours? Keep it. If not, drop it and try another task.

LOW-COST TOOLS

Four tools you can start with today

All four big chat apps have a free version. Pay only when a paid plan saves you real hours.

Sources: openai.com · claude.com/download · gemini.google/overview · x.ai/grok (checked Oct 12, 2026). Check prices on the day.

01 FREE TO START

ChatGPT

By OpenAI. Free to start; paid plans add more use.

02 FREE TO START

Claude

By Anthropic. Free plan; Pro and Max add more use.

03 FREE TO START

Gemini

By Google. The Gemini app is free to start.

04 FREE TO START

Grok

By xAI. Free to try on grok.com and in the app.

OWN YOUR AI

Your data, your AI, your building

Some top AI models are free to download and run on your own computers or a private cloud. You trade some convenience for control.

01

Use an AI service

Easiest. It runs on the vendor's computers, and so does your data.

02

Rent a private computer

Run your own copy of a free model on rented cloud hardware.

03

Run it in your building

Most control. Highest upfront cost, and you maintain it.

MORE CONVENIENT

MORE CONTROL

Source: OpenAI, "Introducing gpt-oss," Aug 2025: the free gpt-oss-120b model (Apache 2.0 license) runs "on a single 80 GB GPU."

BEFORE YOU SWITCH IT ON

Five ground rules

Write them down before the first mistake, not after.

Basis: Trivedi et al., ICML 2026, p. 9 (set rules early); Google AI governance paper, 2026 (existing law still applies); approval settings in OpenAI dots and xAI Grok Bot

TICK ALL FIVE

- ✓ No customer secrets in free tools
- ✓ A person reviews before anything is sent
- ✓ Keep a simple log of what the AI did
- ✓ Tell customers when AI is used
- ✓ Have an off-switch, and know who flips it

DON'T WAIT FOR THE CRASH

Fix it before you're stuck with it

NOW

Write the rules

Ground rules on paper before any tool goes live.

FIRST MONTH

Watch closely

Review every AI output. Keep the log.

FIRST QUARTER

Check the results

Time saved? Mistakes caught? Adjust.

LATER

Harder to change

Once habits set in, fixing them costs more.

"By the time consequences are undeniable, the technology is entrenched and correction is costly."

Source: Trivedi et al., "Solipsistic Superintelligence is Unlikely to be Cooperative," ICML 2026, p. 9 (the Collingridge dilemma)

THINK LIKE A MANAGER

The people who manage AI win

STEP 1

You set the goal

→

STEP 2

AI helpers draft

→

STEP 3

You check quality

→

RESULT ✓

Work you trust

AI 2027 “people who know how to manage and quality-control teams of AIs are making a killing.”

Source: AI Futures Project, “AI 2027” (Apr 2025), PDF p. 8, “Late 2026.” A scenario, not a forecast of fact.

OWNING BEATS EARNING

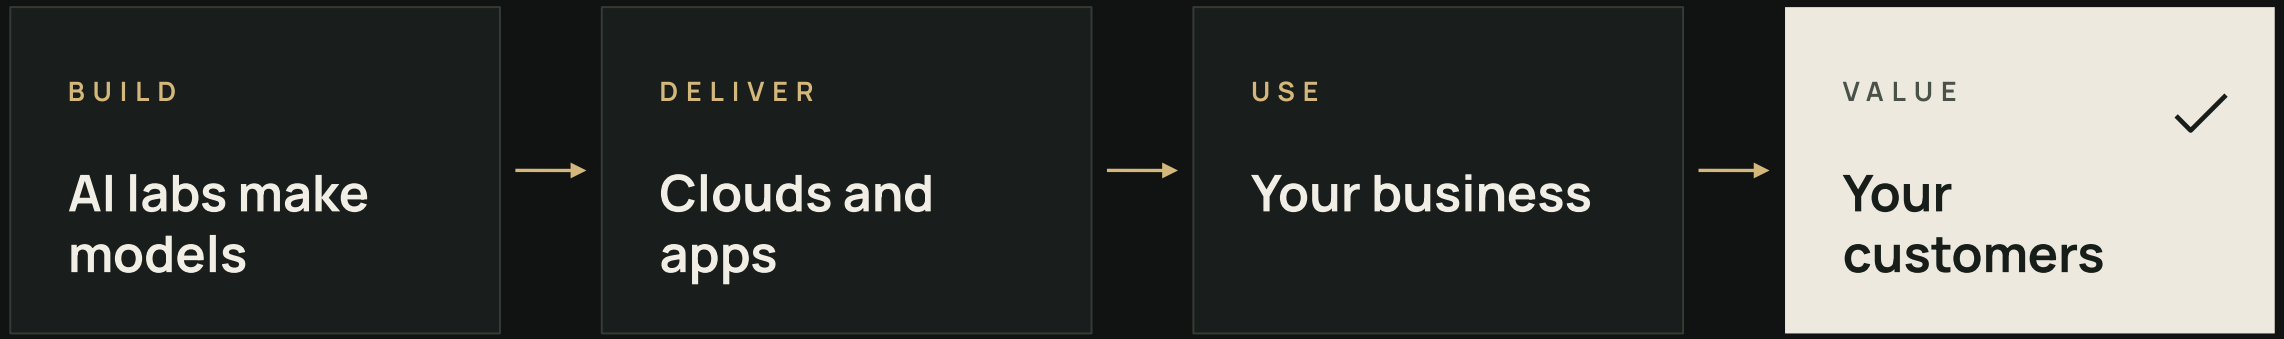
Ownership, not just a check

	A PAYCHECK	A STAKE
What it is	Pay for hours worked	**A share of the gains: an "equity dividend"**
If AI does more of the work	Fewer hours can mean less pay	**You still own your piece**
How the study scores it (if AI fully transforms work)	Job retraining: 5.9 out of 100	**Universal basic capital: 93.5 out of 100**

SOURCE · JACOBS & IMAS, GOOGLE DEEPMIND, "ECONOMIC POLICY FOR AGI," SEP 2026, PP. 4, 14, 18

USE IT WELL

The value lands downstream



LESSON Like electricity, “most of the value accrued downstream to firms using the technology.”

Source: Jacobs & Imas, Google DeepMind, “Economic Policy for AGI,” Sep 2026, p. 19

A PROPOSAL FROM GOOGLE

AI experts inside small businesses

Google proposes short, high-impact “residencies” that put technical experts where AI help is scarce.

Source: Google, “A Pragmatic Approach to AI Governance in America,” 2026, p. 14

01

Small businesses

Hands-on help setting up AI that fits the work.

02

Nonprofits

The same tools, without an IT department.

03

Local governments

Better, faster service for residents.

04

The goal

To “democratize the knowledge” to use enterprise-quality AI.

COMMUNITY FIRST

Tampa builds it together

The fair AI future spreads the gains. AI 2040's plan says "Diffuse AI broadly." We start on Main Street: local, Black-owned, community-first.

Sources: AI Futures Project, "AI 2040: Plan A," summary page; Jacobs & Imas, Google DeepMind, 2026, p. 18



ANTHONY SEALEY · FOUNDER, SOVEREIGN HQ

Tampa, Florida

> *Society has the capacity and tools to shape our economic trajectory in the AGI era.*

— JULIAN JACOBS & ALEX IMAS · GOOGLE DEEPMIND

"Economic Policy for AGI," Sep 2026, p. 26

START WITH ONE USEFUL IMPROVEMENT

Let's talk about *your business.*

sovereign-hq.com ↗

Practical AI for small business · Tampa, FL

SOVEREIGN-HQ.COM

SEALEY.AI