


WELCOME




AI in Practice: A CEO's Guide to Driving Growth and Efficiency

An AI conversation program by the experts at  HatchWorks AI

DOWNLOAD THE SLIDES





Hello, my name is

Brandon Powell

I'm an AI Humanist with a pragmatic lens who helps organizations learn and use AI in ways that actually work in order to transform both their technology and people.



Hello, my name is

Matt Paige

I'm an AI strategist and builder who believes generative AI can empower anyone, anywhere to do amazing things by removing barriers, changing habits, and unlocking human potential.

**Who uses AI on
a daily basis?**





Majority of
organizations are
now ***using*** AI...
(it's now table stakes)

*A new baseline has
been established.*

But there is still a **massive gap** between what is **being done** and what is **possible**

Most organizations are **barely scratching the surface**. The gap between what AI can do and what's actually being adopted is **massive**.




Figure 2: Theoretical capability and observed exposure by occupational category
 This figure shows the share of job tasks that LLMs could theoretically perform (blue area) and our own job coverage measure derived from usage data (red area).

The next **step function** is
redesigning work around **agents**
and adopting a more **agentic**
operating model



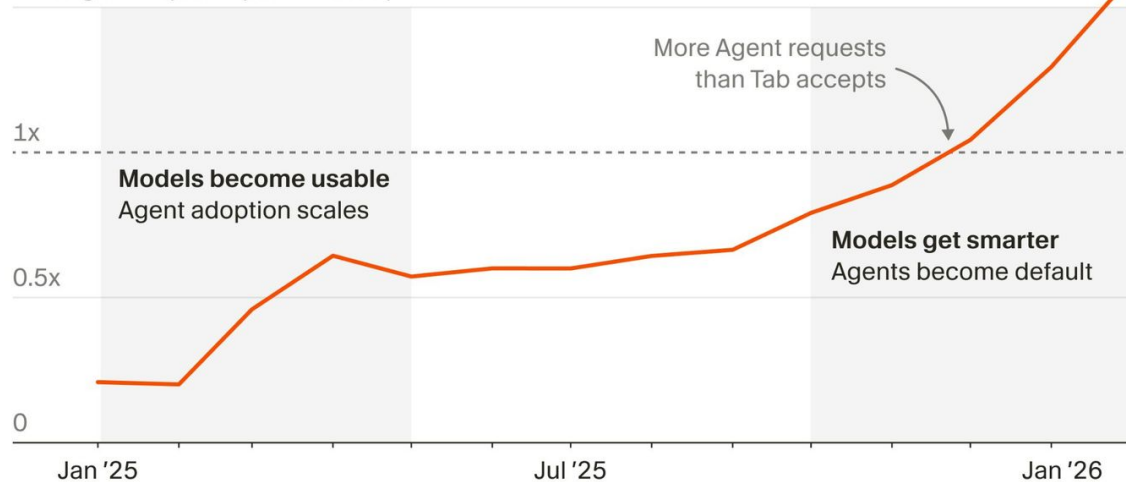


Over the last year, AI has fundamentally shifted from an assistant to an agent.

Agents become mainstream and then the default



1.5x Agent requests per Tab accept



Engineers were first. **Every function is next.**

However...

AI isn't just changing *what* we do.

It is changing *how* organizations
are **built** and **structured**



"CEOs are all arriving at the same question at the same time:
What does an organization look like in an agentic world?"

Mark Moses CEO @  **CEO COACHING[®]**
international

FROM HIERARCHY TO INTELLIGENCE

“Most companies are focused on AI as a productivity enhancer. **Few are focused on the potential of AI to change how we work together.**”

Block is showing what it looks like to fundamentally rethink organization design, ultimately harnessing AI to **increase speed as a compounding competitive advantage.**

Jack Dorsey, CEO @  **BLOCK**



From Hierarchy to Intelligence

Instead of layers of management routing information, they're building what they call a "**company world model**", an AI system that maintains a continuously updated picture of the entire business.



Individual
Contributors



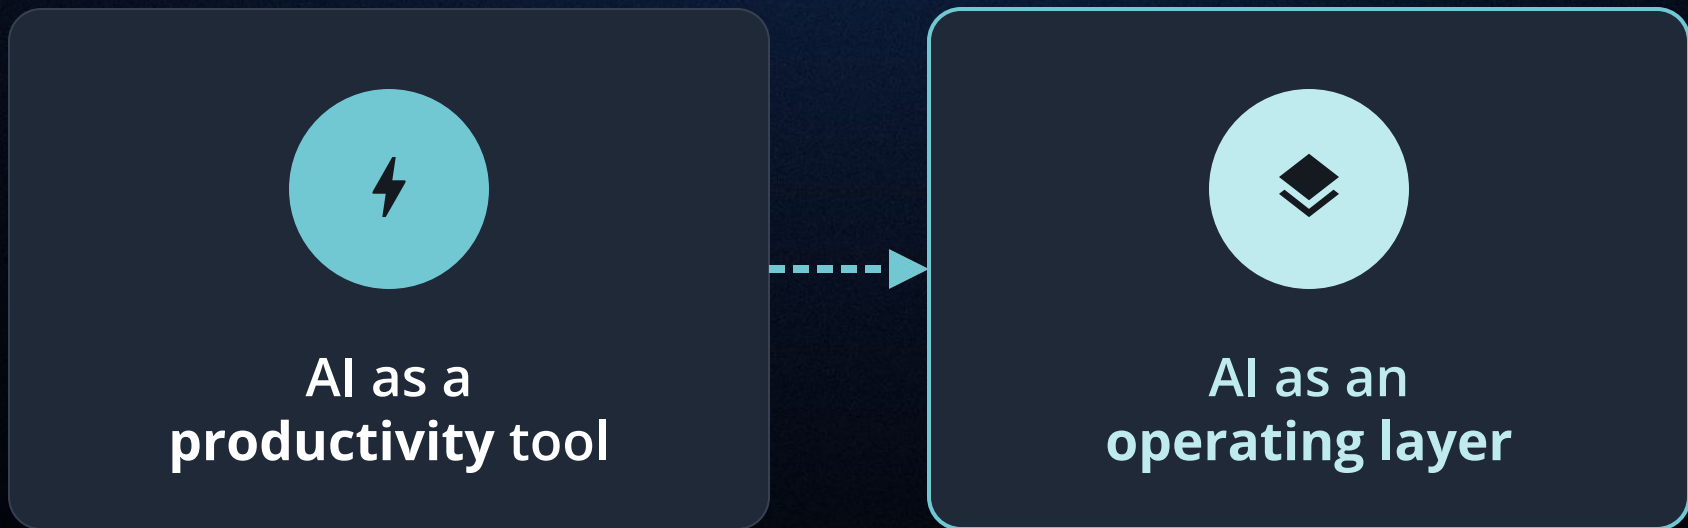
Directly Responsible
Individuals



Player Coaches

~~No middle management~~

Our goal: Help you shift AI from a feature to the foundation your business runs on.



Companies redesigning work (AI Leaders) vs (AI Laggards) layering tools onto old workflows.

• ENTERPRISE

The gap is widening at scale

METRIC	LEADERS	LAGGARDS
Productivity growth	+3.5–5.0%	+0.3–0.8%
EBIT margin impact	+20–30%	Flat
Pilot failure rate	35%	95%

Productivity Divergence (10 years)



[HatchWorksAI

Agents and the Agentic Organization

Software digitized work. **Agents begin to perform it.**

This is the next operational shift: from software that stores information to systems that can execute toward an outcome — '*software that does the work.*'

Software era: captured workflow,
stored records, and improved visibility.

Agent era: plans steps, uses tools,
completes work, hands back the result.

Agents are how AI moves from feature to workforce leverage.

The Emergence of AI Agents


“AI agents are likely
a **multi-trillion dollar
opportunity**... and the
age of AI Agentics is here”

Jensen Huang, CEO @  **nVIDIA**



What is an AI Agent

AI agents are applications that attempt to **achieve a goal** by **observing** the world around them and **acting upon it** using specific **tools** at their disposal.



Model

The LLM or LLMs powering the agent - essentially its brain



Orchestration Layer

Governs reasoning and decision-making



Tools

Enable agents to take action and interact with the world

AI Agents will start by... ✨

1. Automating Routine Tasks



Streamlining repetitive, rules-based activities.

then...

2. Automating Entire Workflows



Orchestrating multi-step sequences and data handover.

then...

3. Automating Entire Job Functions



Executing comprehensive roles and decision-making.

Leading to a shift from task-based automation to role-based execution.

The shift from **Tools** to *Workers*

AI TOOLS (2023-2024)



You ask a question, you get an answer



You're in the loop every step



Saves minutes



Requires you to know what to ask

The Shift

Evolution of Work

AI AGENTS (2025+)



You assign a task, you get a deliverable



The agent plans, executes, and reports back



Saves hours, days, headcount



Requires you to know what outcome you want

This is both **offloading tasks** from you and your team's plate and enabling you to do a whole host of **new ones**.

 Claude

Tactical Agent Playbook

From the experts at  HatchWorks AI

ANTHROPIC

CERTIFIED PARTNER



Claude is one of many harnesses for building agents



NemoClaw



ElevenLabs
Voice AI Agents

HERMES-AGENT



CLAUDE
CODE



Workspace agents
in ChatGPT,

Everything **we will cover...**

What we will cover

Claude Ecosystem

Claude Projects

Claude Connectors

**Scheduled Tasks
& Routines**

Claude Skills

**Claude Agent Strategy
& Managed Agents**

Agents we will show off

Executive Assistant

Sales Agents

Agentic Customer Service

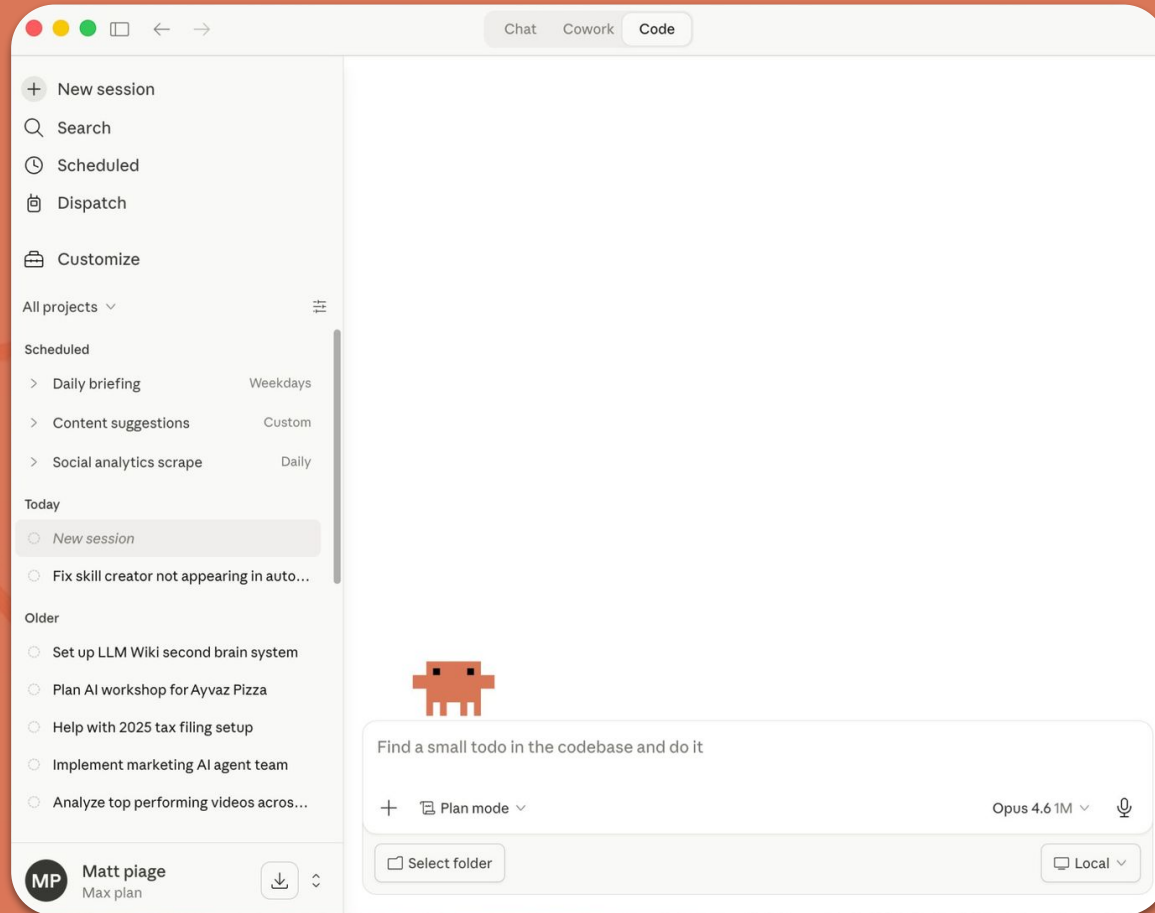
What's the difference between **Claude Chat**, *Claude Cowork*, and Claude Code?

 Claude Chat	 Claude Cowork	 Claude Code
 Your thinking partner	 Your full-time assistant	 Your builder
 Brainstorm, analyze, write, strategize	 Execute multi-step tasks on your files, apps, and systems	 Build software, automations, and agent infrastructure
 Everyone	 Knowledge workers, ops, leadership	 Technical teams, power users
 Conversation	 Assign a task → Claude plans, executes, delivers	 terminal / Claude app / IDEs
 "Analyze this competitive landscape"	 "Review my inbox, calendar, and Slack and build me a morning brief"	 "Build me an internal dashboard that pulls from our CRM API"

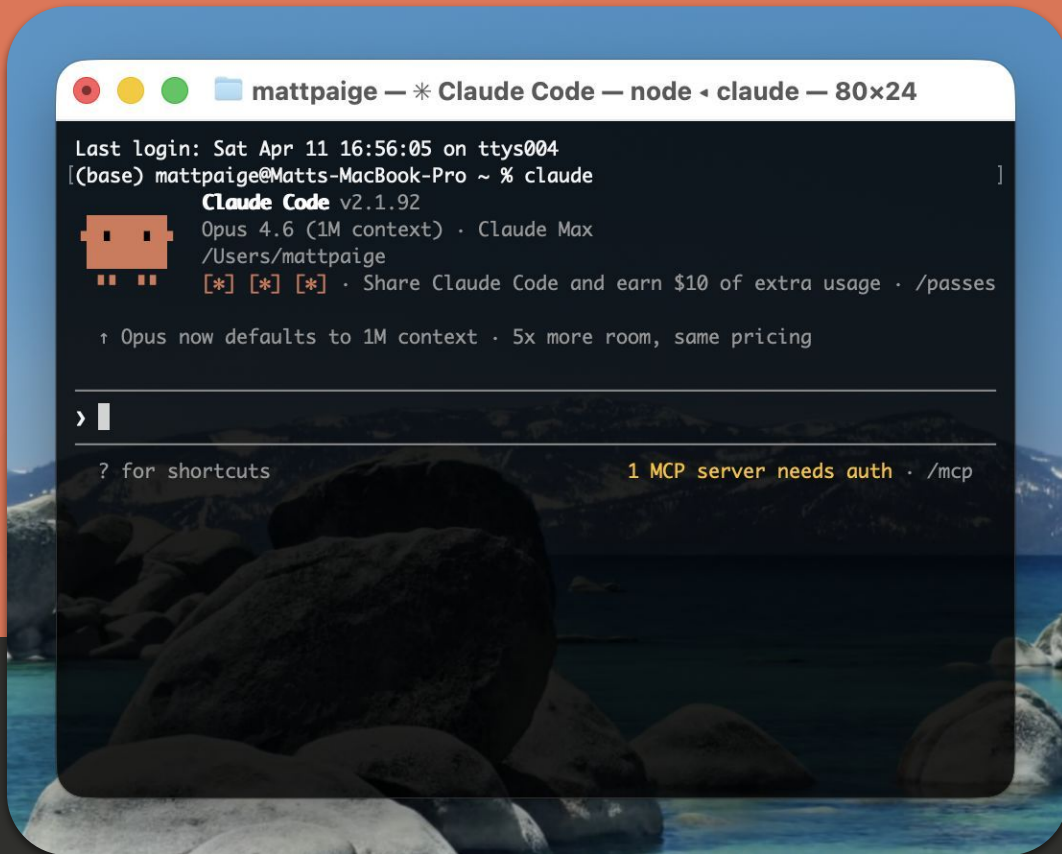
What's the difference between **Claude Chat**, *Claude Cowork*, and Claude Code?

 Claude Chat	 Claude Cowork	 Claude Code	 Open Claw
 Your thinking partner	 Your full-time assistant	 Your builder	 You're always-on agent
 Brainstorm, analyze, write, strategize	 Execute multi-step tasks on your files, apps, and systems	 Build software, automations, and agent infrastructure	 Persistent memory, always on, always learning (open source)
 Everyone	 Knowledge workers, ops, leadership	 Technical teams, power users	 Tinkerers, power users, early adopters
 Conversation	 Assign a task → Claude plans, executes, delivers	 terminal / Claude app / IDEs	 WhatsApp, Telegram, Discord, Slack
 "Analyze this competitive landscape"	 "Review my inbox, calendar, and Slack and build me a morning brief"	 "Build me an internal dashboard that pulls from our CRM API"	 "Watch my inbox and Slack. If anything urgent comes in, ping me on Telegram"

We will be working in the Claude Desktop app. **Why?** Because it is easier.



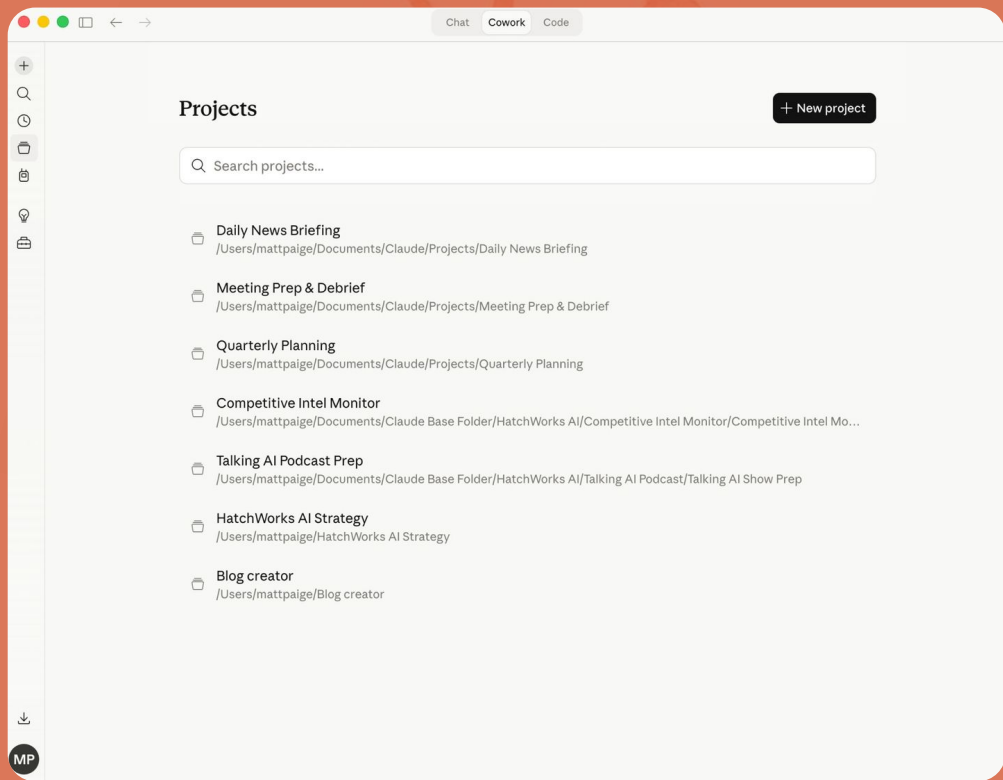
We will be working in the Claude Desktop app. **Why?** Because it is easier.



**No, you don't
need to use the
terminal...**

(unless you want to feel cool)



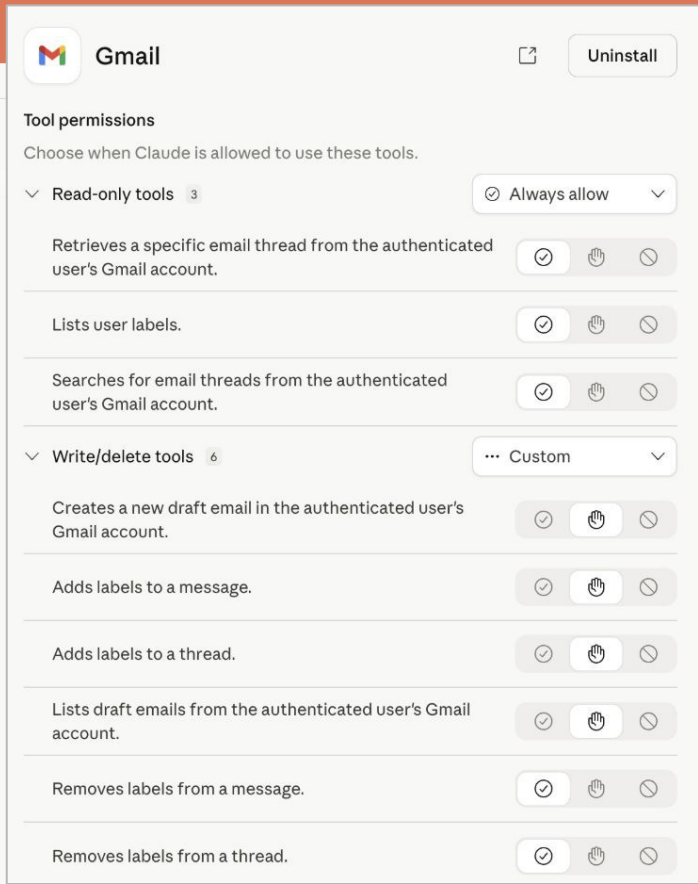
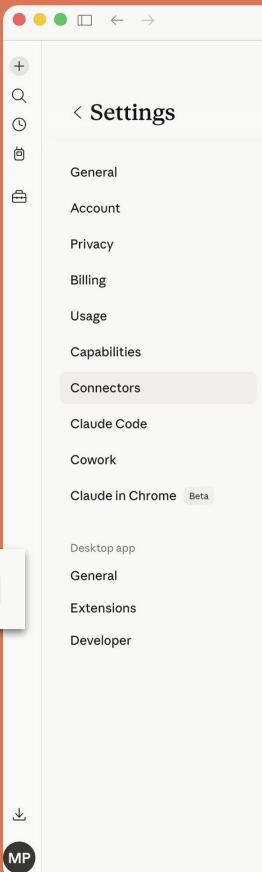
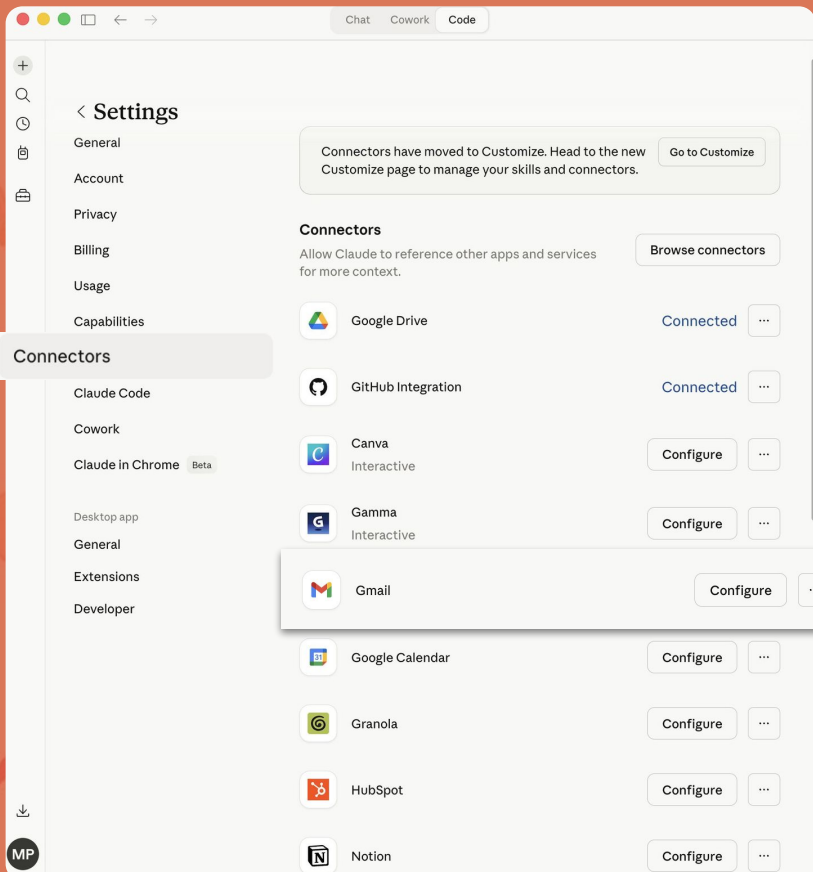


Available in **Chat** & **Cowork**

Organize your agent work into separate workspaces, each with its own files, context, instructions, and memory.

Let's create one.

Claude Connectors



Introducing

salesforce

Salesforce Headless 360

One platform. Anywhere humans & agents work.



First wave of APIs, MCP tools,
and CLI commands **now available**

Build together
with coding agents



Scale your
agents with trust



Deploy on
any surface



Let's build an

Executive Assistant

Providing you a daily brief to start your day



Claude Scheduled Tasks

Scheduled tasks + New task

Run tasks on a schedule or whenever you need them. Type `/schedule` in any existing session to set one up.

Local tasks only run while your computer is awake.

- Daily briefing** Local
Weekdays at ~8:15 AM
Next run: Tomorrow at ~8:15 AM
- Social analytics scrape** Local
Every day at ~6:00 AM
Next run: Tomorrow at ~6:00 AM
- Content suggestions** Local
~At 07:00 AM, only on Monday and Thursday
Next run: Apr 16 at 7:00 AM
- Competitive intel monitor** Local
Every Monday at ~9:00 AM
Next run: Apr 20 at ~9:00 AM

MP Matt piage
Max plan

Available in **Chat & Cowork**

Schedule recurring tasks based on a specified frequency.

Let's schedule one.

CLAUDE CODE ROUTINES

Configure once. Runs on autopilot. Your laptop can be closed.

AI agents that work while you sleep.

Routines go further. They also run the moment something happens that matters.

On a clock

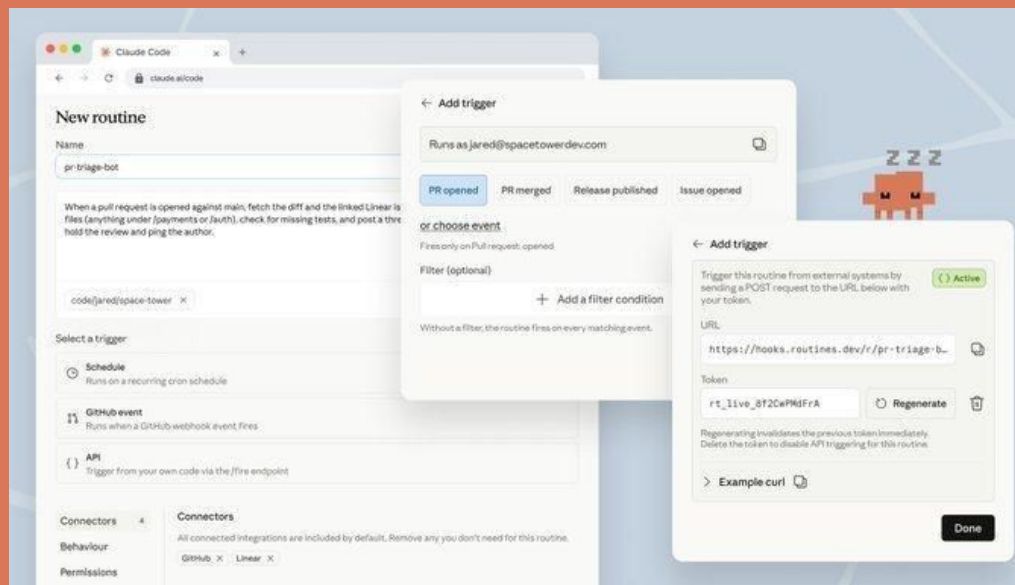
"Run this every morning at 7am." Groom the backlog, audit changes, scan for issues before your team starts the day.

When something changes

"Run this the moment new work is submitted." Every change gets reviewed instantly. No one has to remember to check.


When a system calls for help

"Run this when an alert fires." Your app breaks at 2am. The agent traces the issue and drafts a fix. Your team wakes up to a solution, not a crisis.



Agents are shifting from reactive to proactive

Agents won't be perfect on the first go.



You are the feedback loop.
You got to monitor its
output and help refine it.

P.S. Use Claude to help refine it.

Claude Skills

A markdown file that teaches Claude a repeatable workflow. Write it once, reuse it forever.

```
---
name: your-skill
description: [...]
---

# Your Skill Name

## Instructions

### Step 1: [First Major Step]
Clear explanation of what happens.➡
```



Context-Efficient

Only loads what's needed, when it's needed



Consistent

Same quality output every time, no re-prompting required



Scalable

One skill serves your entire team



skill-name/
└ SKILL.md



1) Frontmatter

Loaded at startup

```
name: check-email
description: check Gmail inbox from past 24
hours, group emails, and draft responses
```

Check Email — 24-Hour Inbox Briefing ## Overview

This skill fetches the user's Gmail inbox from the past 24 hours, organizes threads into logical categories, highlights which emails need a response (with priority ranking), and offers to draft replies using context the user provides.

Instructions

Step 1: Fetch Recent Email Threads

Use the Gmail:search_threads tool to pull all threads from the last 24 hours.

```
Gmail:search_threads
  query: "newer_than:1d"
  pageSize: 50
```

If there are more than 50 threads (a pageToken is returned), paginate to retrieve remaining threads. Continue until all threads from the past 24 hours are collected.

2) Body

Loaded if agent calls skill



I Know Kung-Fu

-NEO

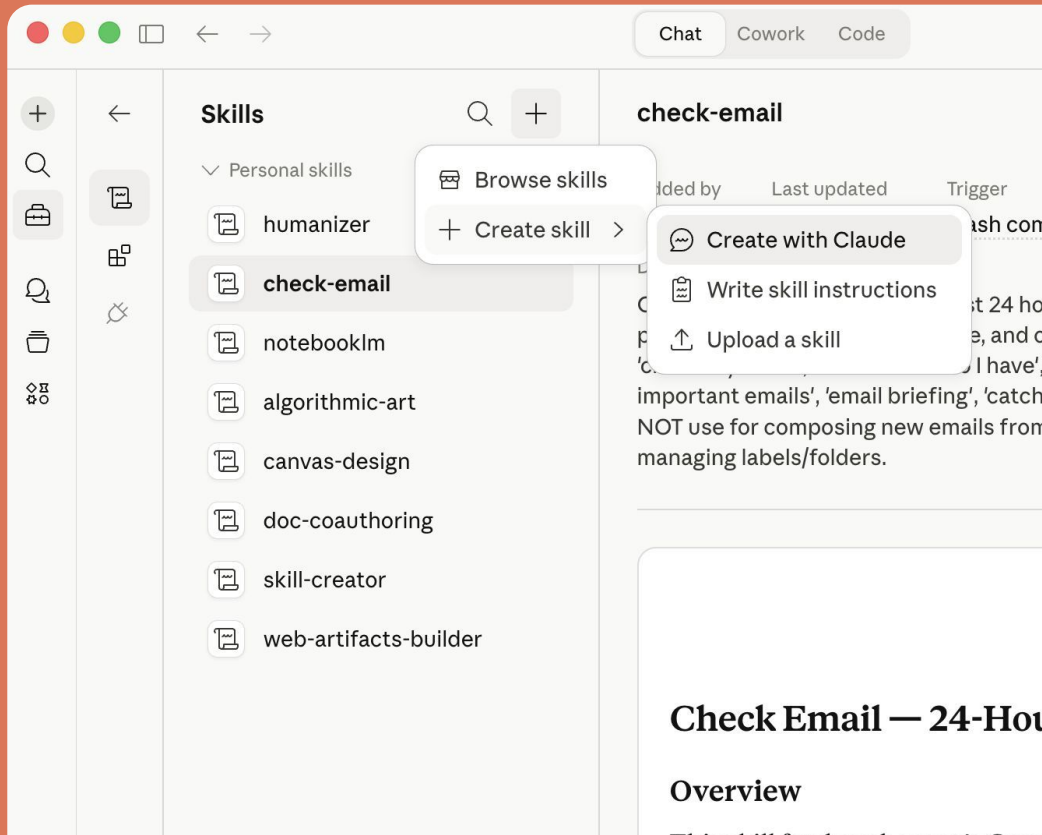
Claude Skills


How do you create a skill?

1. Manually (*Don't do this*)
2. Use the `/skill-creator` skill
3. Upload a skill (GitHub is a great source)

How to use it?


- **/ command**
ie. `/check-email`
- **Related triggers**
(defined in the skill)





SUB-AGENTS

vs



AGENTS TEAMS

Think: *Contractors*

- ▶ Hired for one specific job
- ▶ Work in their own context window, isolated
- ▶ Deliver the result and they're done
- ▶ Don't talk to each other
- ▶ Best for independent, parallel tasks
- ▶ Lower token cost

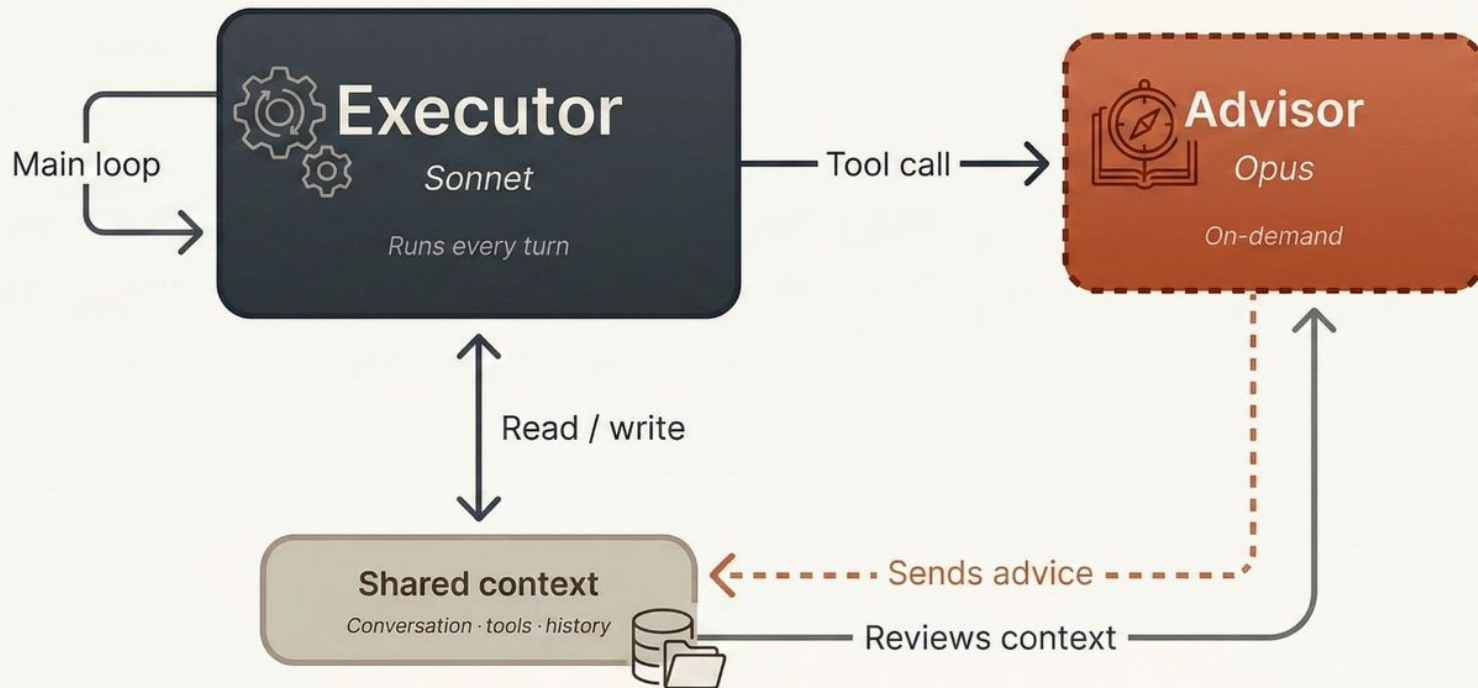
Isolated work → Sub-Agents

Think: *Colleagues*

- ▶ Hired to work together long-term
- ▶ Lead + teammates + shared task list
- ▶ Message each other in real time
- ▶ Build shared context over time
- ▶ Higher token cost, more power
- ▶ Best for collaborative work

Collaborative work → Agent Teams

The Advisor Strategy



Advisor reads the same context as Executor

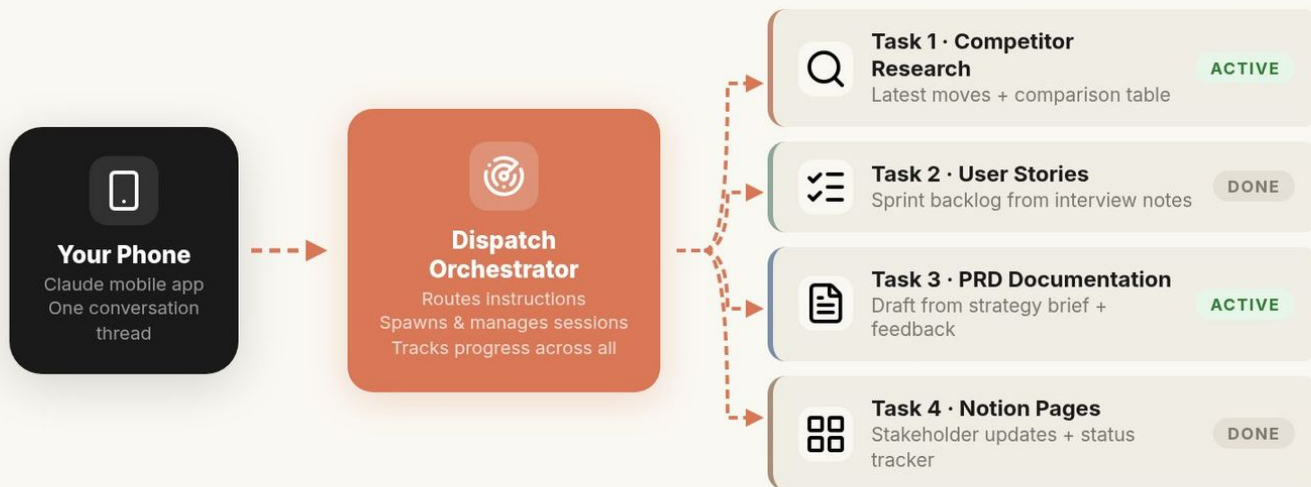
Let's build a

Sales Intelligence Agent



Your Phone Is the **Command Chair**

Dispatch is an orchestrator, not a chat app. One conversation on your phone spawns and manages multiple parallel task sessions running on your desktop.



Your Desktop · Claude Cowork

Files, connectors, CLAUDE.md — all running while you're away

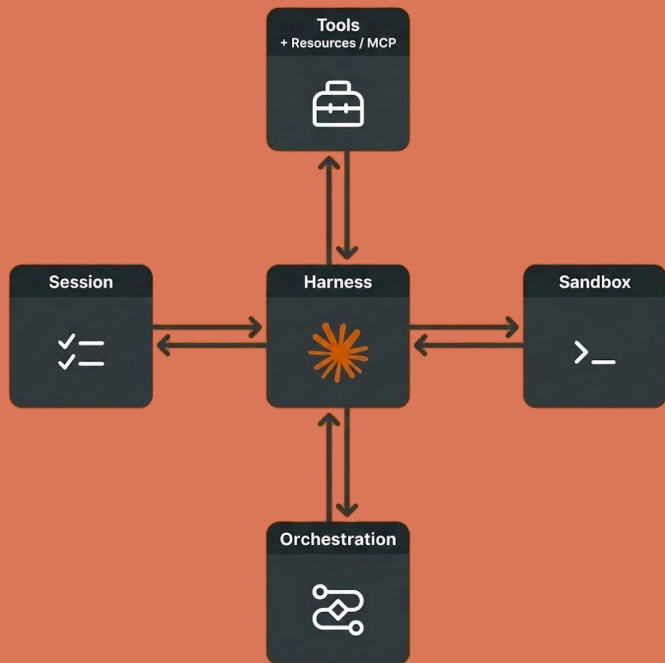
ANYWHERE

CLOUD

YOUR MACHINE

CLAUDE MANAGED AGENTS

The infrastructure to deploy AI agents at scale is here. **The question is how you use it.**



Harness

The brain. Manages how the agent thinks, plans, and recovers when things go wrong.

Tools

What the agent can access. Your CRM, databases, documents, APIs.

Sandbox

Secure workspace. Runs code and handles files without touching production systems.

Session

Persistent memory. Tasks run for hours, survive interruptions, pick up where they left off.

Orchestration

Multiple agents working together. One agent assigns and manages the others.

The platform is ready. The strategy is the hard part.

Defining the right agents, scoping permissions, connecting systems, ensuring governance. That's where the real work lives.

Let's see a

Biz Dev Prospectin g Agent



Now let's see a
**Full Customer
Service
Agent Team**
with Voice in Action



The starting line has been redrawn. You have hired your first agent.

The question is no longer whether to do this... it's where do I start?



[:] HatchWorksAI

PS. This QR code gives you the step by step guide to recreate all the agents we just walked through.

An AI conversation program by the experts at [:] HatchWorksAI



Six Agentic Strategic Principles



Re-imagine Processes

Don't redesign the legacy process around AI



Onboard Agents

Treat agents like you are onboarding employees



Use Cases with Clear ROI

Pick use cases with real ROI measurable day 1



Purpose-built Agents

"Don't build one agent that does everything."



The RIGHT Clean Data

You only need the data required for a use case.



Behavioral Shift

Treat like a transformation and not like a tool

Agentic Strategic Principle #1: Reimagine, Don't Automate

THE OLD WAY: AUTOMATE



Don't just automate your existing processes.

THE NEW PRINCIPLE: REIMAGINE



Step back, and reimagine how they should work now that you have AI + agents.

Agentic Strategic Principle #2:

Onboard Agents Like Employees

You wouldn't hire someone and say "*go figure it out.*"
Don't do that with agents either.



Agentic Strategic Principle #3:

The Agent Team Model

Don't build one agent that does everything. Build purpose-built agents that each own a function.



Briefing Agent

Information synthesis & prioritization



Pipeline Agent

Revenue intelligence & deal tracking



Research Agent

Market & competitor monitoring



Onboarding Agent

Customer setup & onboarding communication

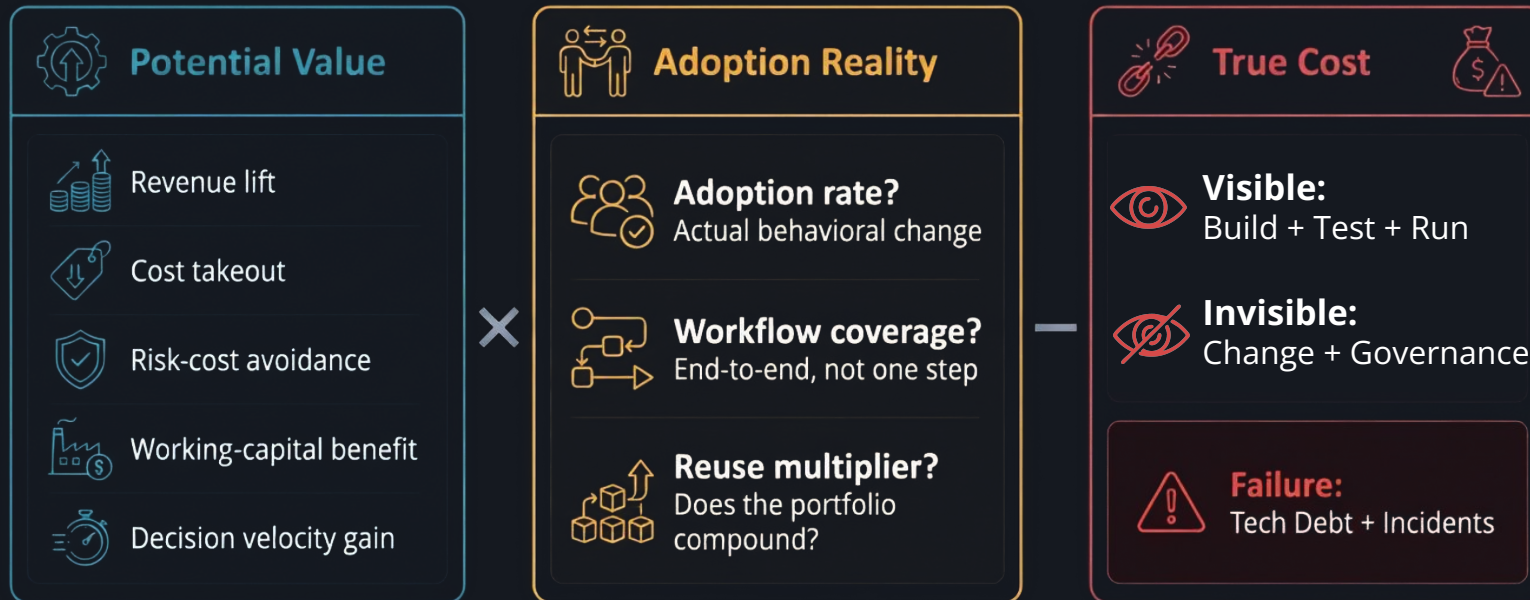


Reporting Agent

Weekly & monthly metric compilation for entire company

Agentic Strategic Principle #4: Pick Use Cases with Clear ROI

Pick a use case with real ROI measurable from day 1



Most business cases only **budget the first two** (Potential Value & Adoption Reality). They **ignore True Cost**.

Agentic Strategic Principle #5: **You Only Need the Right Data**

Data is the new oil and it's your moat but you only need the data required for the use case to be clean



Your proprietary data (context) is your moat

An agent with proprietary business knowledge provides a decisive advantage over generic ones.



Target the right data to assess, validate & store

Identify, clean, and govern the specific data your agent needs. Avoid over-scoped data projects and focus on shipping.

ACCESS ≠ ADOPTION



Agentic Strategic Principle #6:
Treat AI Like a Behavior Shift

And not like a tool tool that you give your teams.

HOW TO THINK ABOUT IT

Every process fits one of three tiers.

The tier determines how — and whether — to deploy an agent.

**IF A HUMAN
FOLLOWS A
CHECKLIST, AN
AGENT CAN OWN IT.**



TIER 1: AUTOMATE

Rules-based. High-volume.
Agent executes.

- ✓ Invoice processing
- ✓ Data entry
- ✓ Report generation

**IF THEY SPEND 80%
GATHERING AND
20% DECIDING,
AUGMENT THE 80%.**



TIER 2: AUGMENT

Agent prepares.
Human decides.

- 🔍 Deal research
- 🔍 Proposal drafting
- 🔍 Code review

**IF YOU'RE ASKING
"HOW DO WE SPEED
THIS UP?" — ASK "WHY
DO WE DO IT THIS WAY?"**



TIER 3: REIMAGINE

Process changes.
Agent enables.

- 🚀 3-person teams replacing 12
- 🚀 Continuous compliance monitoring
- 🚀 Real-time decision loops

YOUR MONDAY MORNING

Three questions before your first production agent.

01



Where's the biggest friction in my business?

Not the shiniest AI use case. The process where your team loses the most time, makes the most errors, or creates the most bottlenecks. That's where your first agent lives.

RULE OF THUMB: If your team follows a checklist → **AUTOMATE IT**.
If they spend 80% gathering, 20% deciding → **AUGMENT THE 80%**.

02



Can an agent actually do this job today?

Five things determine readiness:

- Is the process documented?
- Is the data clean?
- Is there a human review model?
- Can your team collaborate in real time?
- Is governance in place?

Most failures aren't AI failures — they're documentation and collaboration failures.

03



Will my people adopt and can I keep score?

The ROI formula:


$\text{Value} \times \text{Adoption} - \text{Cost}$


Start with one purpose-built agent and one KPI, then decide in 90 days to scale, revise, or stop.

Access ≠ Adoption.
A treadmill in every home doesn't cure heart disease.


Agentic Opportunity **FINDER**


Find your first agent. 3 minutes. On your phone.

 **Where's the friction?** *Which process costs you the most time and money?*

 **Is it documented?** *Could a new hire follow the steps without shadowing someone?*

 **Is the data ready?** *Can you access what you need without manual exports?*

 **Will your team adopt?** *When you roll out new tech, do people actually use it?*

 **What's it worth?** *Can you estimate the cost of this bottleneck today?*

You'll get a named agent recommendation,
a readiness score, and a clear next step.



SCAN TO FIND YOUR FIRST AI AGENT

9 questions · 3 minutes
Personalized recommendation

[:] HatchWorks AI

IF YOU NEED HELP

Three ways we help after today.



01

BUILD YOUR AI ROADMAP

Develop your company AI strategy and roadmap to give you a blueprint tailored to your business.

For leaders who need the map.



02

LAUNCH YOUR FIRST AGENT

Take your first high value use case and make it a reality by building and launching in production agents

For teams ready to ship, not pilot.




03

GAIN COMPETITIVE ADVANTAGE

We help you leverage your data as the AI foundation to increase profitability and differentiation.

For organizations going all-in.




**“This is the worst
AI will ever be...”**



Michael Luo, Stripe

We are just at the
beginning of this
disruption...

Talking AI NEWSLETTER



CUT
THROUGH
THE AI HYPE

Subscribe for practical
AI tips and strategies



PODCAST

www.HatchWorks.com/talking-ai





Get the *slides*

[:] HatchWorksAI

AI in Practice: A CEO's Guide to Driving Growth and Efficiency

An AI conversation program by the experts at [:] HatchWorksAI

