November 3, 2026 | Hosted by Endsight

AI Office Hours

endsight



Jason Clause

Director of Marketing

- 20+ years in managed IT services
- Bachelor's in Public Relations
- Microsoft Certified Professional



Brian Tirado

Director, AI & Automation

- 25+ years in technology
- Leads Endsight's AI & Automation practice
- Guides clients through AI platform selection, rollout, integration, and adoption



Michelle Brezenski

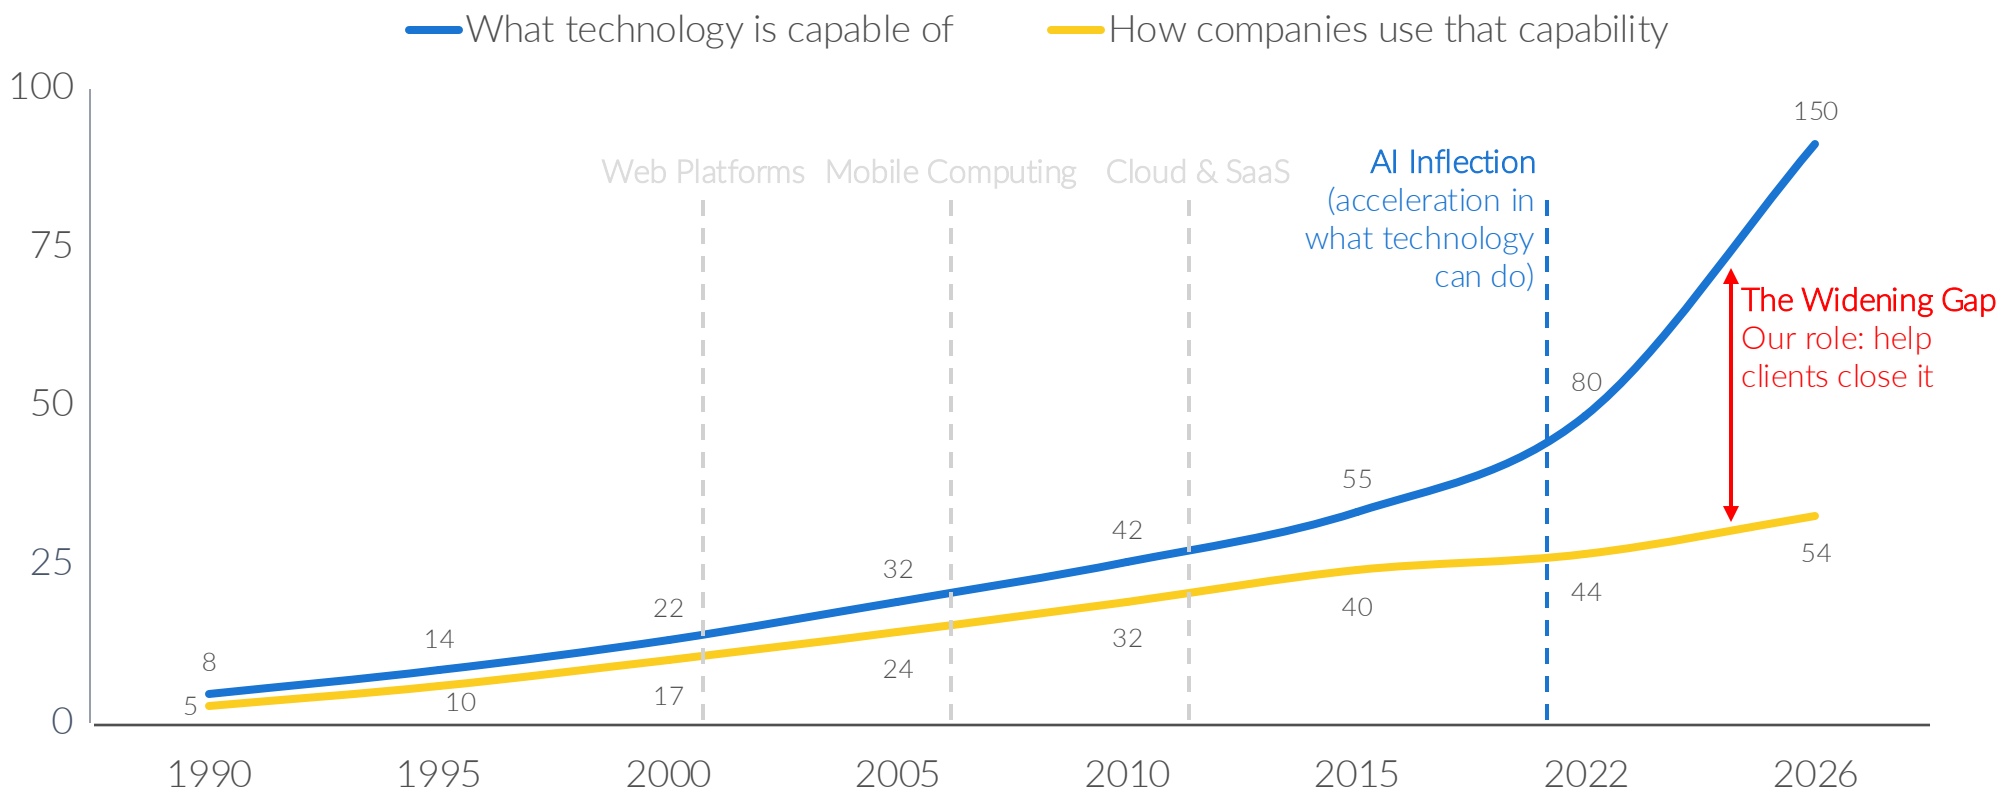
Manager, Development

- 18 years at Endsight
- Multiple Microsoft Certifications
- Led Endsight's development department since 2014

Today's Agenda

1. Welcome and introductions – 3 minutes – Jason
2. The AI landscape: 90-day view – 12 minutes – Brian
3. Client case example: Q&A with David Smith – 20 minutes – All
4. Next steps and resources – 5 minutes – Jason
5. Open Q&A – 20 minutes – All

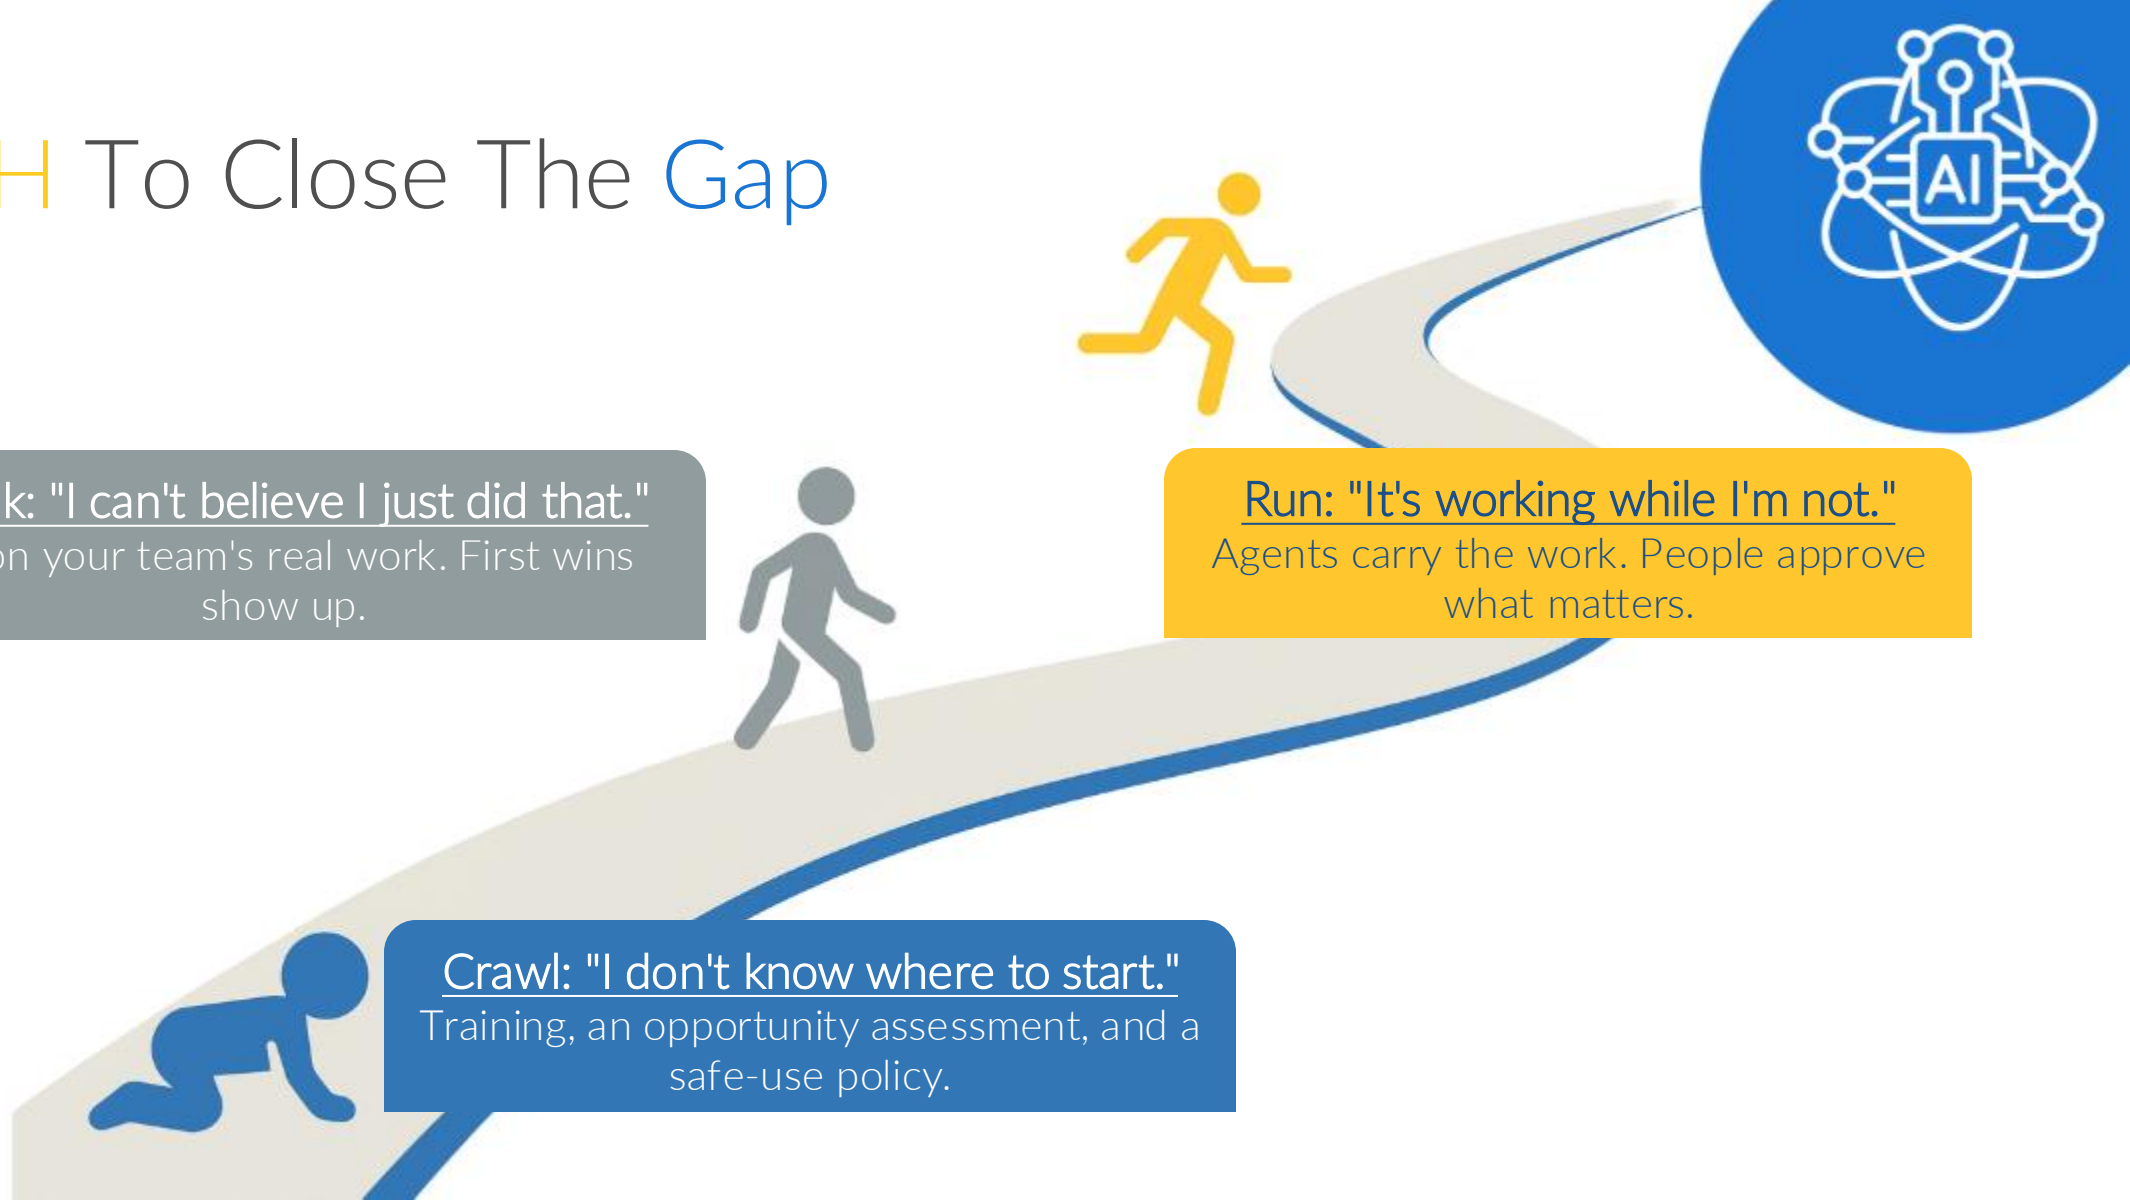
The Technology **Capability** Gap



- Technology **has always outpaced** how we use it
- AI has made the distance between what is possible and what most companies use **dramatically larger**
- The gap is turning into a chasm

endsight

A PATH To Close The Gap



Run: "It's working while I'm not."
Agents carry the work. People approve what matters.

Walk: "I can't believe I just did that."
AI on your team's real work. First wins show up.

Crawl: "I don't know where to start."
Training, an opportunity assessment, and a safe-use policy.

Where Teams Get **Stuck** With AI

Licenses nobody opens
You bought seats. Most sit untouched after month one.

One power user
One person uses AI for everything. Everyone else waits.

Tool sprawl
ChatGPT here, Copilot there. Nobody knows what's in use.

AI that can't see your business
Generic answers, because AI can't reach your files or systems.

Unclear what's safe
Nobody's sure what's OK to paste in, or where it goes.

Pilots that never ship
A promising demo, a few weeks of buzz, then nothing in production.

Almost every one of these comes down to **people.**

The Three Kinds of AI **Users**



The Innovator

Grabs the tool and experiments without being told to. A deployment strategy points them at the problems that matter most.



The Curious but Stuck

Knows there's something here but doesn't know where to start. One good first win on their own work is usually all it takes to unlock them.



The Holdout

Has decided this isn't for them yet. They rarely get argued into it. They come around when they watch the people next to them start winning.

How important is A.I. strategically to the business?

endsight

Your AI Strategy: Three Focus **Streams**



Fluency

Everyone can use AI well, safely, on their real work.

Where you feel it

Licenses nobody opens. One power user, everyone else waiting.

Start this year

About five hours of basics for everyone, then hands-on work by team.



Data

AI can reach the information your people already use.

Where you feel it

Generic answers, because AI can't reach your files or systems.

Start this year

Pick one company-controlled AI tool and connect what you already use.



Security

Protection that keeps pace with how fast AI moves.

Where you feel it

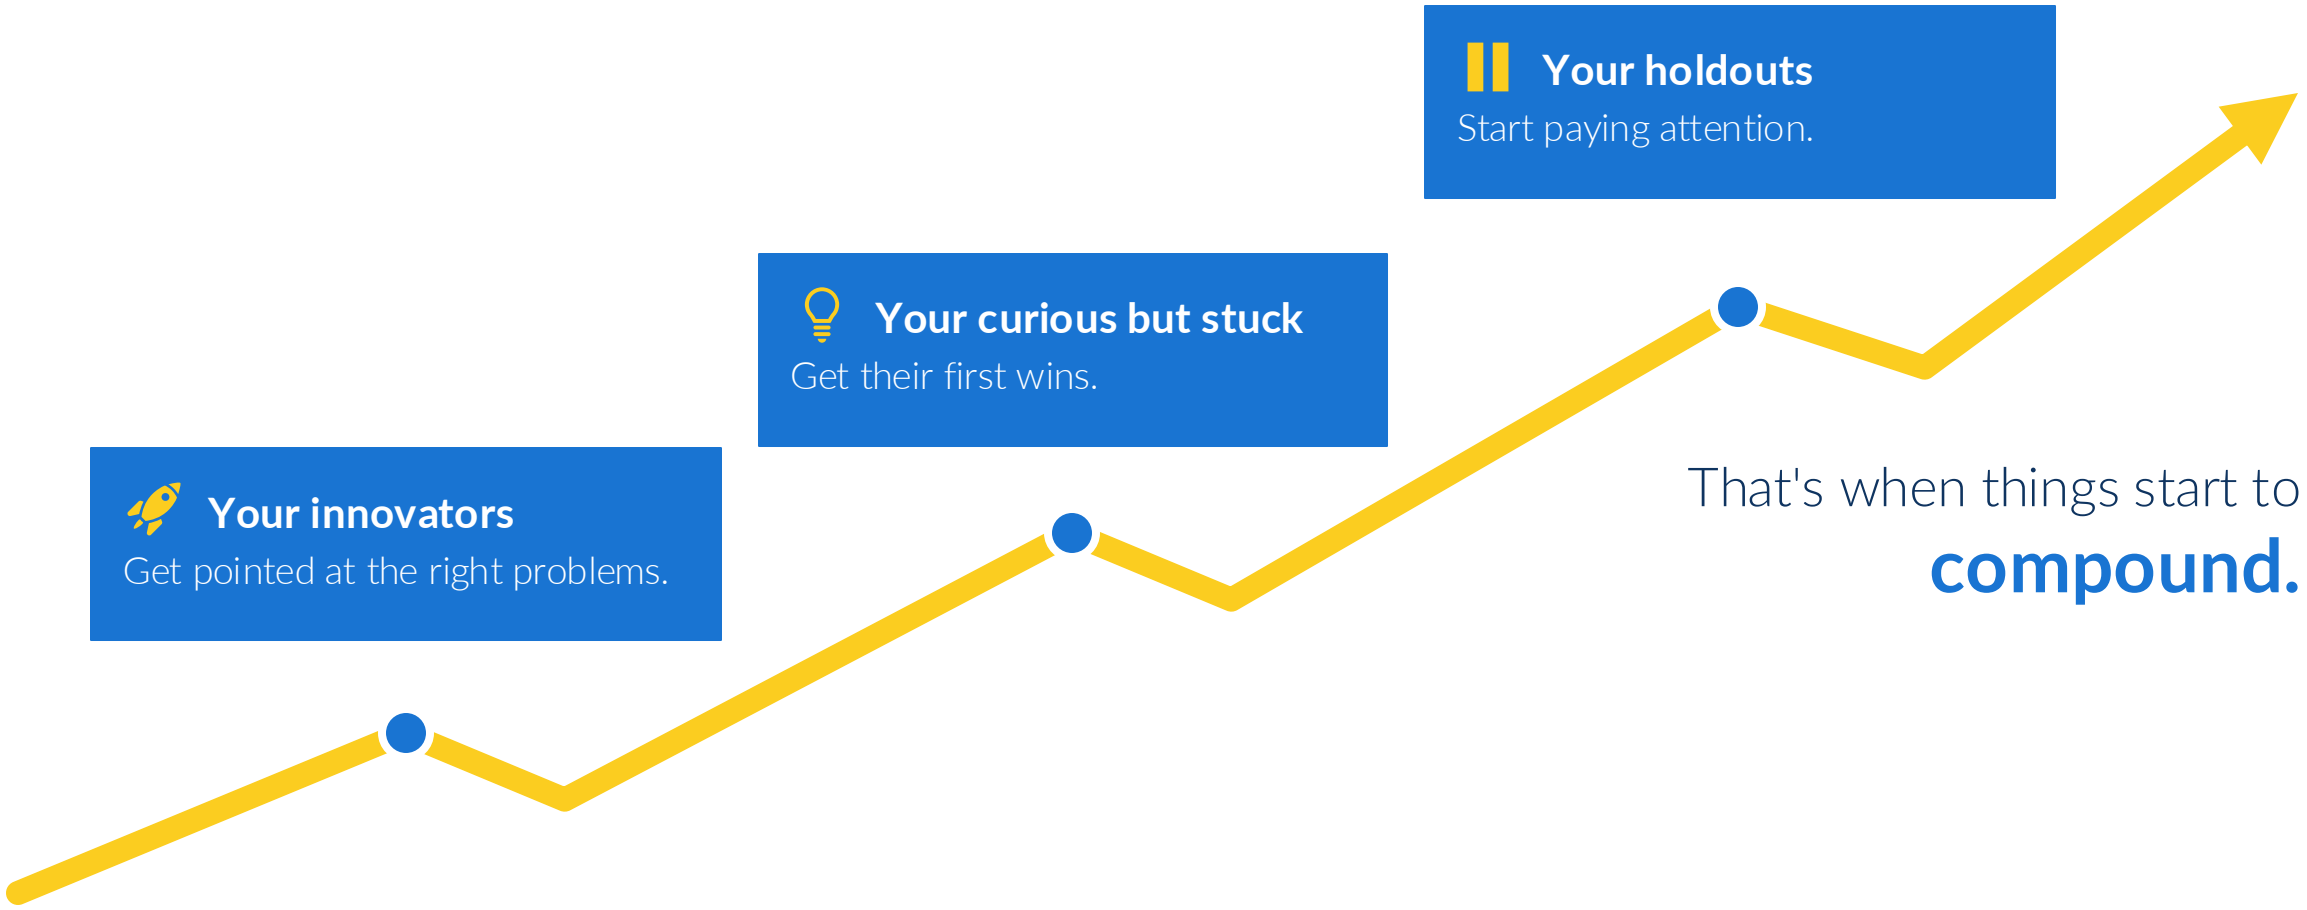
Nobody's sure what's safe to paste in, or where it goes.

Start this year

A written safe-use policy, and a list of every AI app connected to company email and files.

Run all three **at once**, over the next three to five years. Fluency without data stalls, and data without security is a risk you didn't mean to take.

What Happens On The **Path**



BRIAN TIRADO

The AI Landscape: A 90-Day View

endsight

ENDSIGHT.NET

From Drafts to DISCOVERY

Last quarter AI started finishing the work. This quarter it started solving problems experts have worked on for decades.

722

original math papers released by OpenAI in one drop on October 6, from about 4,000 AI attempts.

88 hours

for 10,000 AI agents working together to solve a Navier-Stokes Millennium Prize problem. The Clay Institute says it appears settled.

4 vendors

SpaceXAI, Meta, OpenAI, and Anthropic each shipped a personal agent that works in the background.

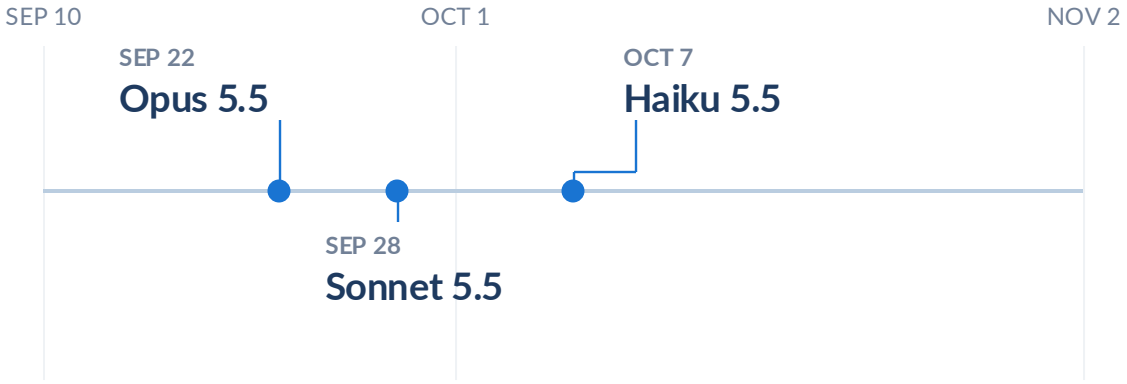
Why it matters

The models solving open math problems are the same kind your team uses for proposals, analysis, and operations. Capability is moving faster than most business plans for it. The playbook: one governed, business-tier platform, with a named person accountable for anything that leaves the building.

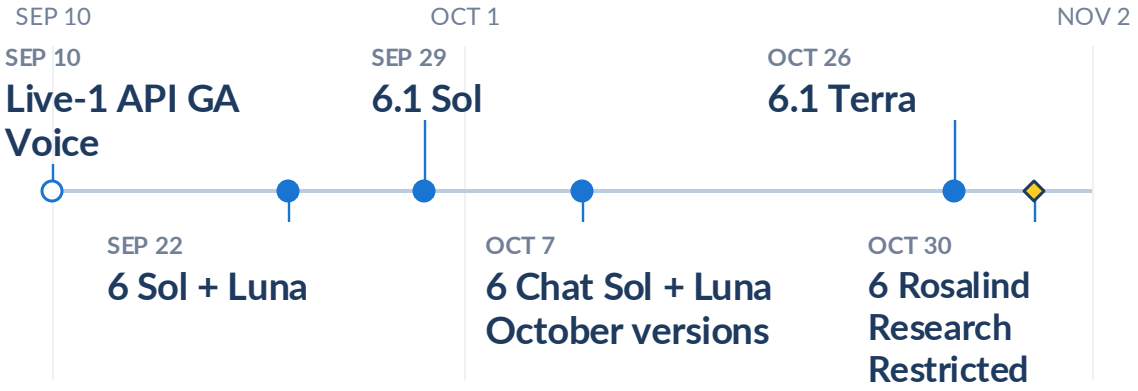
AI model RELEASES & access

Vendor timelines · Same date scale in every panel · September 10–November 2, 2026

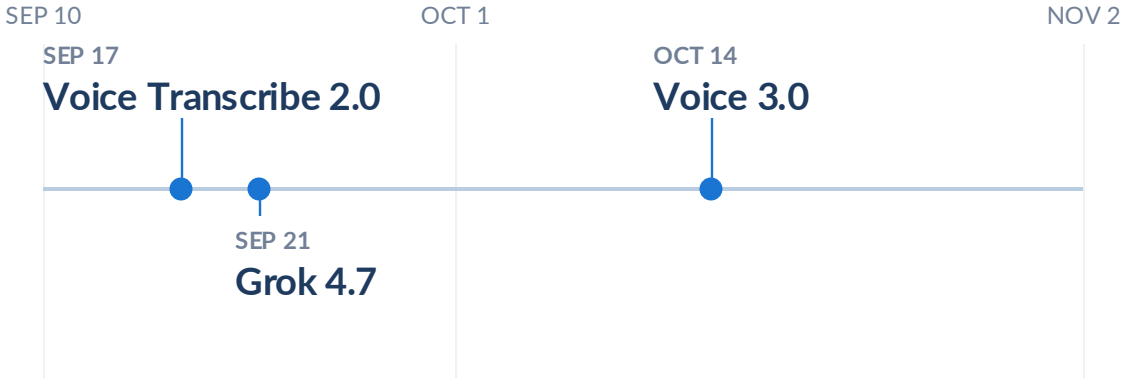
Anthropic / Claude



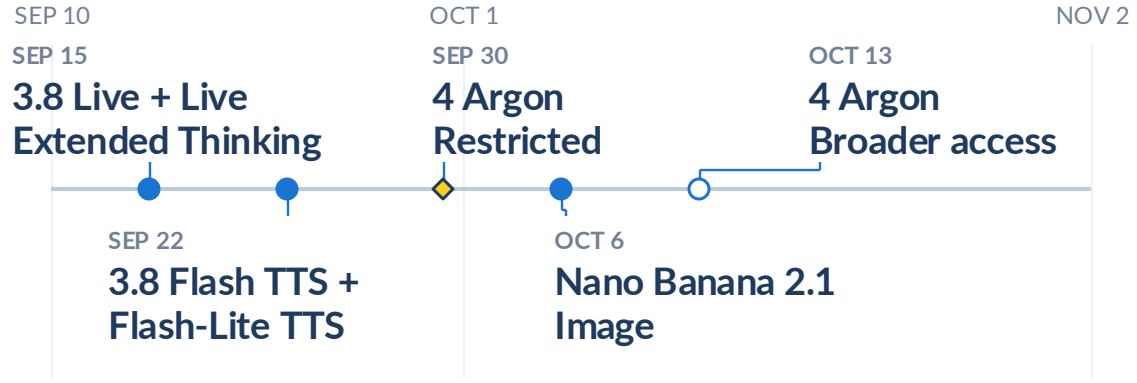
OpenAI / GPT



SpaceX / xAI / Grok



Google / Gemini



endsight

New model / version · Restricted release · Expanded access

AI reasoning PROGRESS

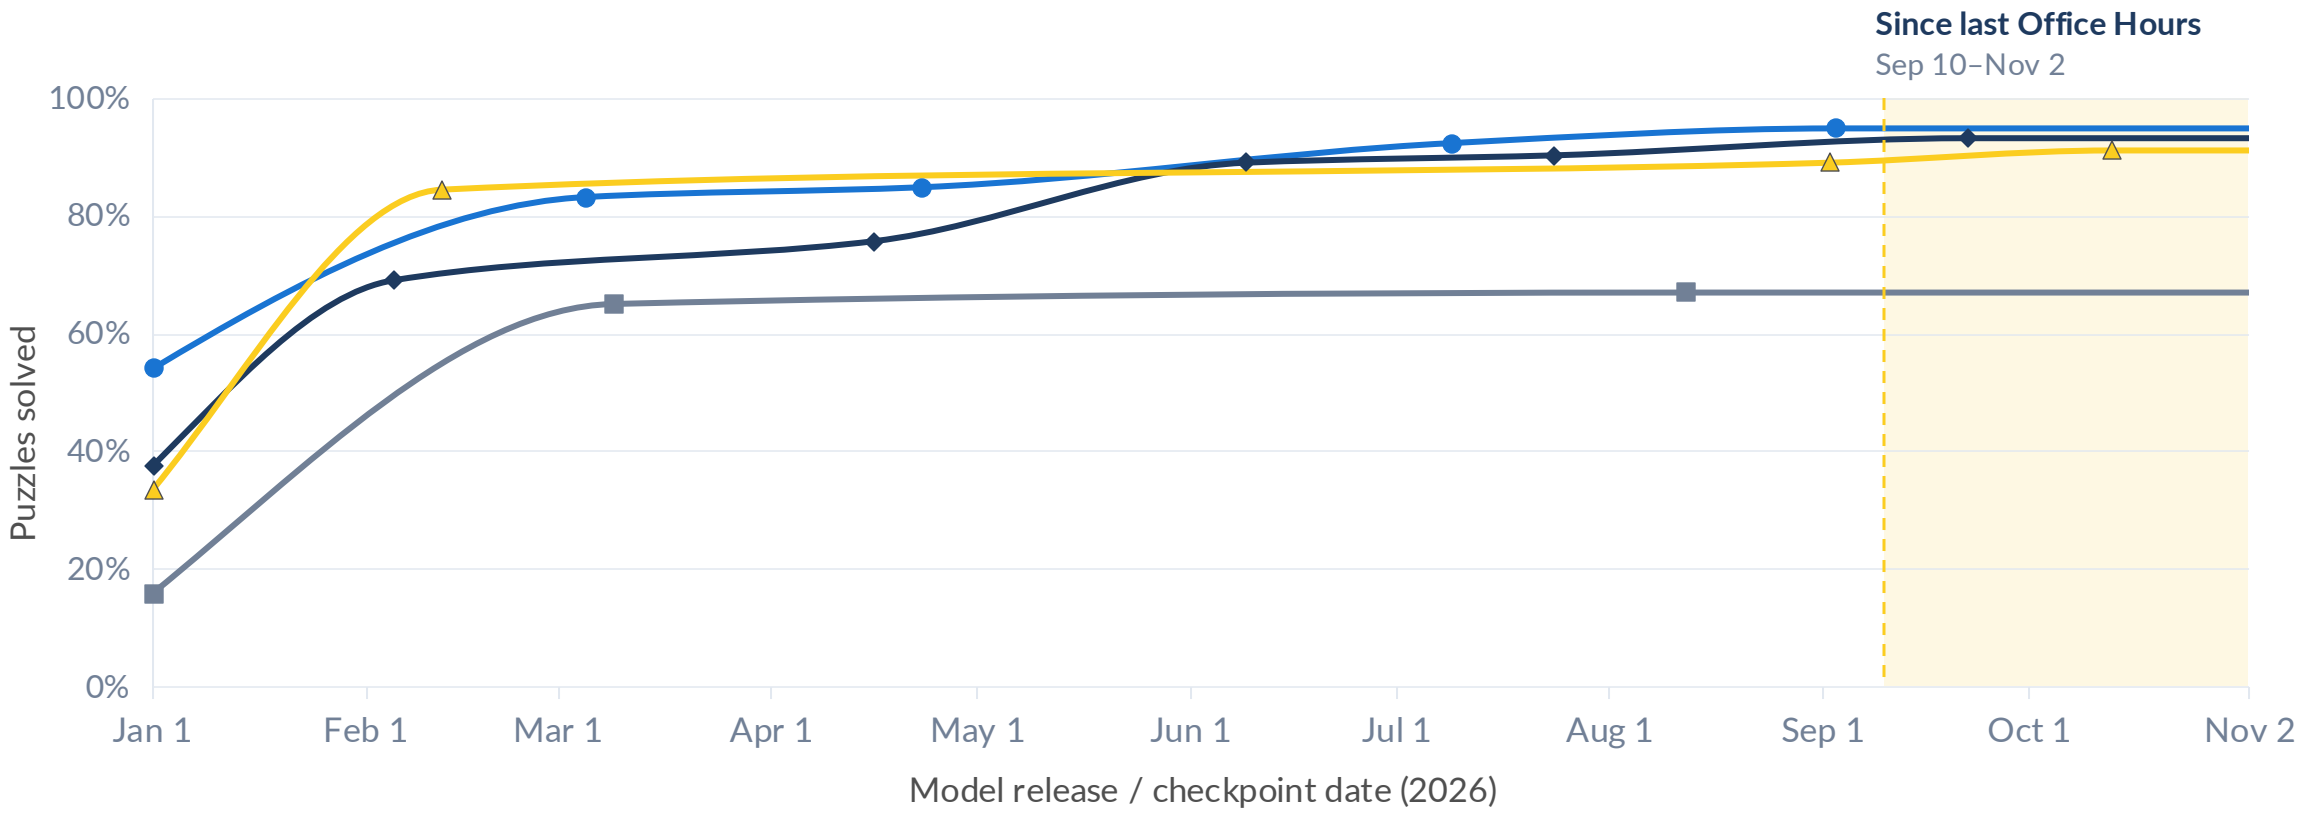
Highest reported ARC-AGI-2 puzzle scores by vendor · January 1–November 2, 2026

OpenAI · 95.0%
GPT-6 Astra · Sep 3

Anthropic · 93.3%
Claude Opus 5.5 · Sep 22

Google · 91.3%
Gemini 4 Argon · Oct 13

xAI · 67.1%
Grok 4.6 · Aug 12



endsight

ENDSIGHT.NET

Source: ARC Prize semi-private leaderboard, Oct 28 snapshot. November 2 endpoints are carried forward; full sources and dates in notes.

AI Did New MATH

Why a math story belongs in a business briefing: it is the clearest sign yet that AI can do original work, not just summarize existing work.

Navier-Stokes

A $1M question about fluids

The equations behind weather forecasts, wing design, and blood flow. The open question: can the math itself break down?

OpenAI says 10,000 agents working 88 hours showed it can, when an outside force is applied. The no-force version is still open.

OpenAI, Sep 8, 2026

The 722 drop

Research at scale

722 papers from about 4,000 attempts, averaging about three hours of AI time per accepted result.

That is research at a pace no human team can match.

OpenAI, Oct 6, 2026

Original work, not just drafts. The same kind of model is on your team's desktop.

Personal Agents ARRIVED

In ten weeks, four major AI companies shipped an agent that works in the background and asks before it acts.

What they all do

- **Watch your email and calendar** and bring you what needs attention
- **Read Teams or Slack** and find loose ends
- **Prep you** for upcoming meetings
- **Work in the background** on each vendor's most capable model
- **Use a computer or browser** as a tool to finish tasks
- **Ask first** before paying, signing, sending, or deleting

OpenAI Dots	**Always on, in chat** Background agents on GPT-6 Astra that answer in Slack and Teams. Pro and Business Premium plans.	*OpenAI, Sep 29, 2026*
Claude Duet	**Watches, then asks** Morning brief from email and calendar; keeps working with your laptop closed. Research preview for US Pro and Max; not yet on Team or Enterprise.	*Anthropic, Oct 21, 2026*
Meta Muse	**Watch it work, step by step** Visible browser on a cloud computer, with free, $20, and $100 tiers. Meta can train on your queries unless you opt out.	*Meta, Sep 8, 2026*
Grok Bot	**Learns by watching you** Show it a process once and it saves a routine. Bundled with SuperGrok and Cursor plans.	*SpaceXAI, Aug 11, 2026*

Same shape, four vendors. Each one signs into your accounts and asks before sensitive steps.

The capability is HERE

A certified partner, certified architects, and a help desk already working AI tickets.

Certified

Claude Partner

Endsight is a Claude Certified Partner in the Claude Partner Network: trained and vetted to design, deploy, and support Claude for business.

10

Claude Certified Architects

Ten Claude Certified Architects on staff, the credential Anthropic issues to partners who meet its design and deployment standard.

Hundreds

AI tickets worked last quarter

Our Response Center has worked hundreds of AI-related tickets over the last quarter. Day-to-day support for AI tools is here, and the capability keeps growing.

What to DO

Where we would start, and a new capability we can bring to it.

Our platform recommendation

Claude as the primary platform. The strongest fit for business work, and the one we support end to end.

Sonnet 5.5 day to day, Opus 5.5 for the hardest work. Most tasks run on Sonnet; reserve Opus for complex analysis and agent builds.

Business-tier plans with a shared skill library. Company-managed accounts, SSO, and repeatable skills the whole team can use.

Anthropic Defender program

Endsight has been accepted. A program for partners doing security work using Anthropic's platform.

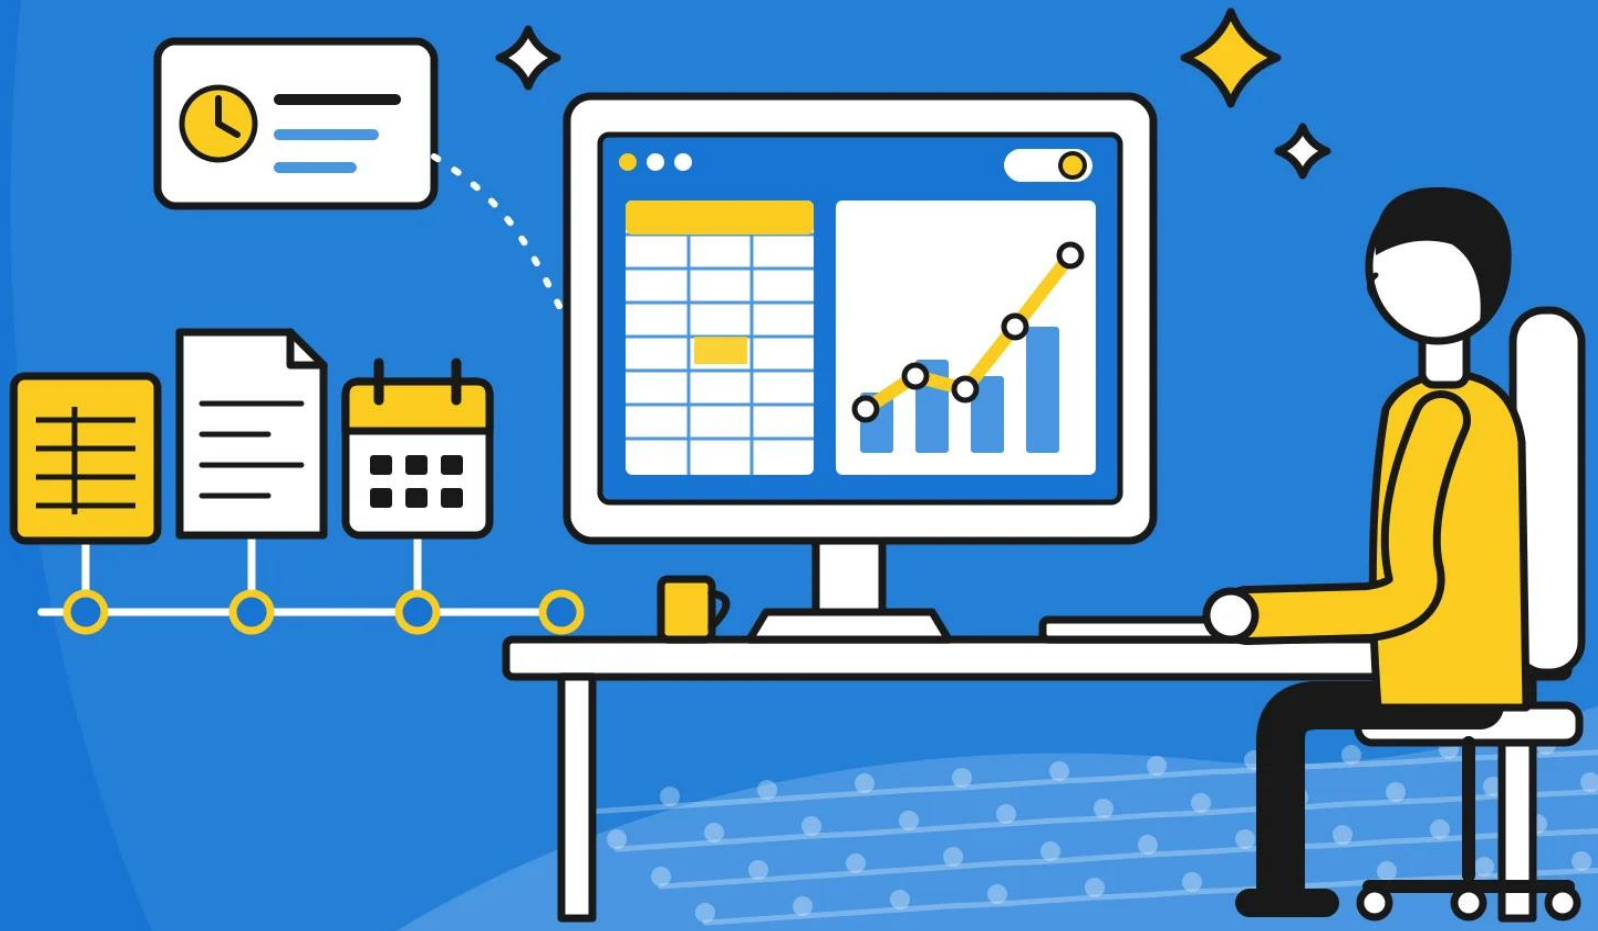
Stronger application reviews. Builds on our existing capability to review agent-built and vibe-coded applications.

Access to non-public models. Models not available to the public, used to find weaknesses before an attacker does.

endsight

CUSTOMER SUCCESS STORY

DMS Strategic Financial Advisory Reclaims 800+ Hours a Year with Claude



What David BUILT

Figures from the Endsight case study for DMS Strategic Financial Advisory, published September 30, 2026.

800+

hours a year reclaimed across two recurring workflows alone

16 hrs

to build the 2027 annual operating plan, a job that normally takes three months

75 to 3

minutes for the daily cash reconciliation, about 438 hours a year returned

"I'm not a computer programmer. But I now am an Excel computer programmer, because I ask it in English to do something I know can be done in Excel, and it does it in a minute."

David Smith, Founder and Fractional CFO, DMS Strategic Financial Advisory

What you'll hear from David

- How he works with AI in plain English, no coding required
- The planning, forecasting, and cash work you do every month
- What changed, and what it took to get there

Read The Full **Case Study**

How a fractional CFO practice stopped experimenting with AI and started running it

1

The challenge

Planning and reconciliation rebuilt by hand, every time, from exports out of several systems.

2

The approach

Claude working inside Excel, with each recurring analysis saved as a reusable skill he runs on demand.

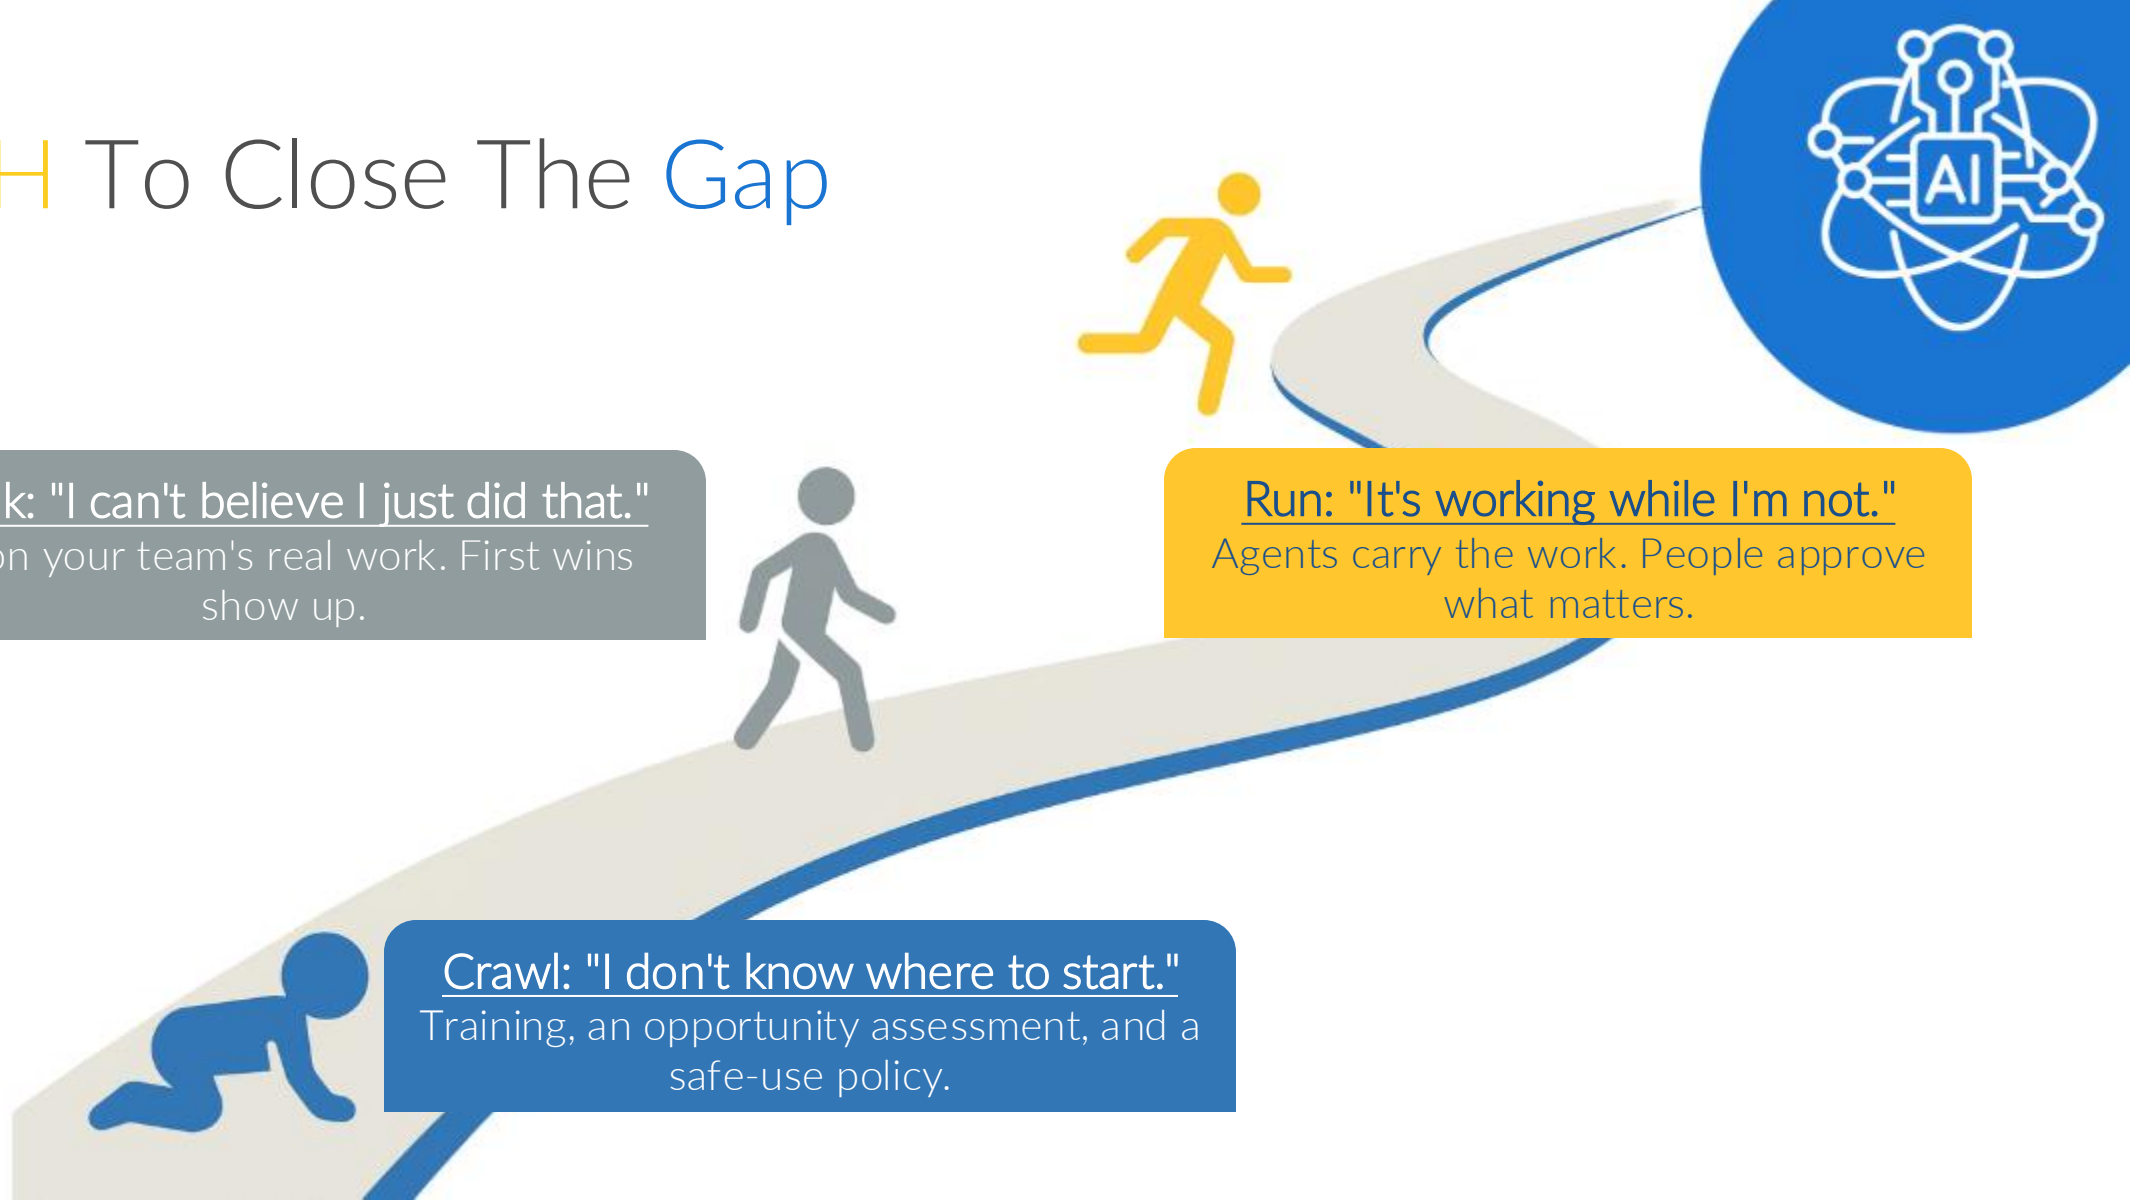
3

Why it worked

Each skill captures David's own judgment, so every later run takes minutes instead of days.

Scan to read the full story

A PATH To Close The Gap



Run: "It's working while I'm not."
Agents carry the work. People approve what matters.

Walk: "I can't believe I just did that."
AI on your team's real work. First wins show up.

Crawl: "I don't know where to start."
Training, an opportunity assessment, and a safe-use policy.

Your PATH, Your DOORS, Your CHOICE.



Keep Tuning In

Still putting the pieces together? That's okay. We'll keep bringing you the best AI content we can. Please keep tuning joining us.

Want to Talk It Through

Some clients just want a conversation. Call us. No agenda, no pitch, we'll meet you wherever you are.

Ready for a Structured First Step

The AI Deployment & Training Bundle is a guided, structured way to get real work done with AI. It's a custom path built around your team, your tools, and your pace.

Building Something & Need a Sherpa

If you're already moving but want an experienced practitioner alongside you, we can help you build.

Something else Resonated

If today sparked an idea that doesn't fit neatly into any of these, we want to hear it. There's no wrong door.

Want Help BUILDING?

Business judgment

How your work gets done

- **Workflows:** how work really moves today
- **Cost:** what each step and error costs
- **People:** who does the work, who has to change
- **Risk:** what can't go wrong
- **Control:** where a person signs off

+

Technical judgment

What it all runs on

- **Models:** which AI, or plain automation, fits
- **Data:** where it lives and how clean it is
- **Systems:** Microsoft 365, SharePoint, your CRM
- **Security:** AI sees only what each person can
- **Reliability:** monitoring, plus a help desk

Put both on one team, and the way your business works becomes **work AI can do.**

Rules-based automation

Inputs and rules are predictable

An AI agent

The goal is clear, the path varies

A person in control

The decision carries real risk

We Hope You'll Join Us Again Soon.

Cybersecurity End-User Training

December 8th

Practical, engaging security training for your whole team, not just IT. Approved for MCLE credit.



Cybersecurity Office Hours

January 28th

A quarterly session with Stephen Hicks covering the latest threats, what they mean for your organization, and what to do next. Open to clients and the broader business community.



AI Office Hours

February 25th

Our next session is coming up February 25th, we hope to see you there!



We'd love to hear from you. A short survey is on its way. Your feedback shapes what we cover next.

First Principles Podcast Launches This NOVEMBER!

FIRST PRINCIPLES

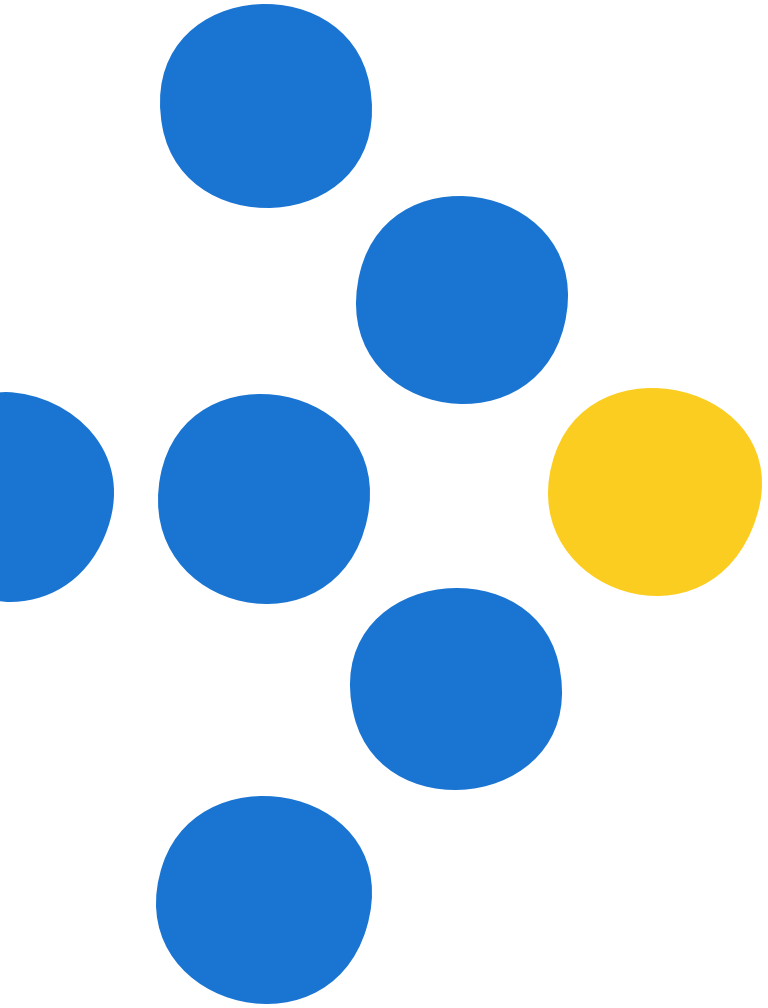
Where proven operators debate business in the open.

Part podcast, part community. Mike Chaput and fellow operators put leadership ideas on the table and test them against real-world experience.



LAUNCHING SOON!

Scan to subscribe



Questions

Pre-Submitted QUESTIONS

Law firm finance

Submitted before the session

For a mid-sized law firm's finance function, where do you **trust Claude to handle the heavy lifting** on things like client profitability or hourly "productive" tracking, and where do you **still require a human review** before anything goes to partners?

Thank You

endsight

Most Common Questions From Client Meetings

Is our data safe? Will the AI train on what we put in? — Asked in ~18 of 26 calls

How much training do our people need, and what does it look like? — Asked in ~16 of 26 calls

What does it cost? How is it licensed? — Asked in ~14 of 26 calls

Which platform should we pick — Claude, ChatGPT, or Copilot? — Asked in ~13 of 26 calls

Do we need an AI acceptable use policy, and what goes in it? — Asked in ~13 of 26 calls

Can you connect the AI to our line-of-business systems? What's an MCP? — Asked in ~13 of 26 calls

What does deployment actually involve? Why not just buy licenses ourselves? — Asked in ~12 of 26 calls

People are already using AI on their own — how do we get shadow AI under control? — Asked in ~10 of 26 calls

Will AI expose data people shouldn't see? — Asked in ~10 of 26 calls

How does the Microsoft 365 connector work — is it mirroring our data? — Asked in ~10 of 26 calls

What happens when an employee leaves? Who owns the account and the data? — Asked in ~10 of 26 calls

Where do we start? Which use case first? — Asked in ~10 of 26 calls

Team plan or Enterprise plan? — Asked in ~9 of 26 calls

How do usage limits work? Will we blow through them? — Asked in ~9 of 26 calls

Will AI replace our people? How do we bring fearful staff along? — Asked in ~9 of 26 calls

How do we standardize and share what works — skills, prompts, our voice? — Asked in ~9 of 26 calls

What happens after deployment and training — how do we keep getting value? — Asked in ~9 of 26 calls

Can we trust the output? What about hallucinations? — Asked in ~8 of 26 calls

Should we pilot or roll out to everyone? — Asked in ~7 of 26 calls

Our systems don't talk to each other — what about our data? — Asked in ~7 of 26 calls

Security & Data Privacy

- Is our data safe? Will the AI train on what we put in?
- Will AI expose data people shouldn't see?
- What happens when an employee leaves? Who owns the account and the data?

endsight

Can we trust the output? What about hallucinations?

endsight

Cost & Licensing

- What does it cost? How is it licensed?
- Team plan or Enterprise plan?
- How do usage limits work? Will we blow through them?

endsight

Platform & Technical Integration

- Which platform should we pick — Claude, ChatGPT, or Copilot?
- Can you connect the AI to our line-of-business systems? What's an MCP?
- How does the Microsoft 365 connector work — is it mirroring our data?
- Our systems don't talk to each other — what about our data?

endsight

Governance & Policy

- Do we need an AI acceptable use policy, and what goes in it?
- People are already using AI on their own — how do we get shadow AI under control?

endsight

Getting Started & Rollout

- Where do we start? Which use case first?
- Should we pilot or roll out to everyone?
- What does deployment actually involve? Why not just buy licenses ourselves?

endsight

Training & Change Management

- How much training do our people need, and what does it look like?
- Will AI replace our people? How do we bring fearful staff along?
- How do we standardize and share what works ? (skills, prompts, MCPs)

endsight

What happens after deployment and training? How do we keep getting value?

endsight