Spetember 10th, 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 - Ten minutes. Brian
3. Demo: Smile Back => Google Review Request – Fifteen minutes Jason
4. Model Context Protocol (MCP) - Ten minutes. Brian & Michelle
5. Next steps and resources – 2 Minuets
6. Open Q&A – 20 Minutes.

The Technology **Capability** Gap

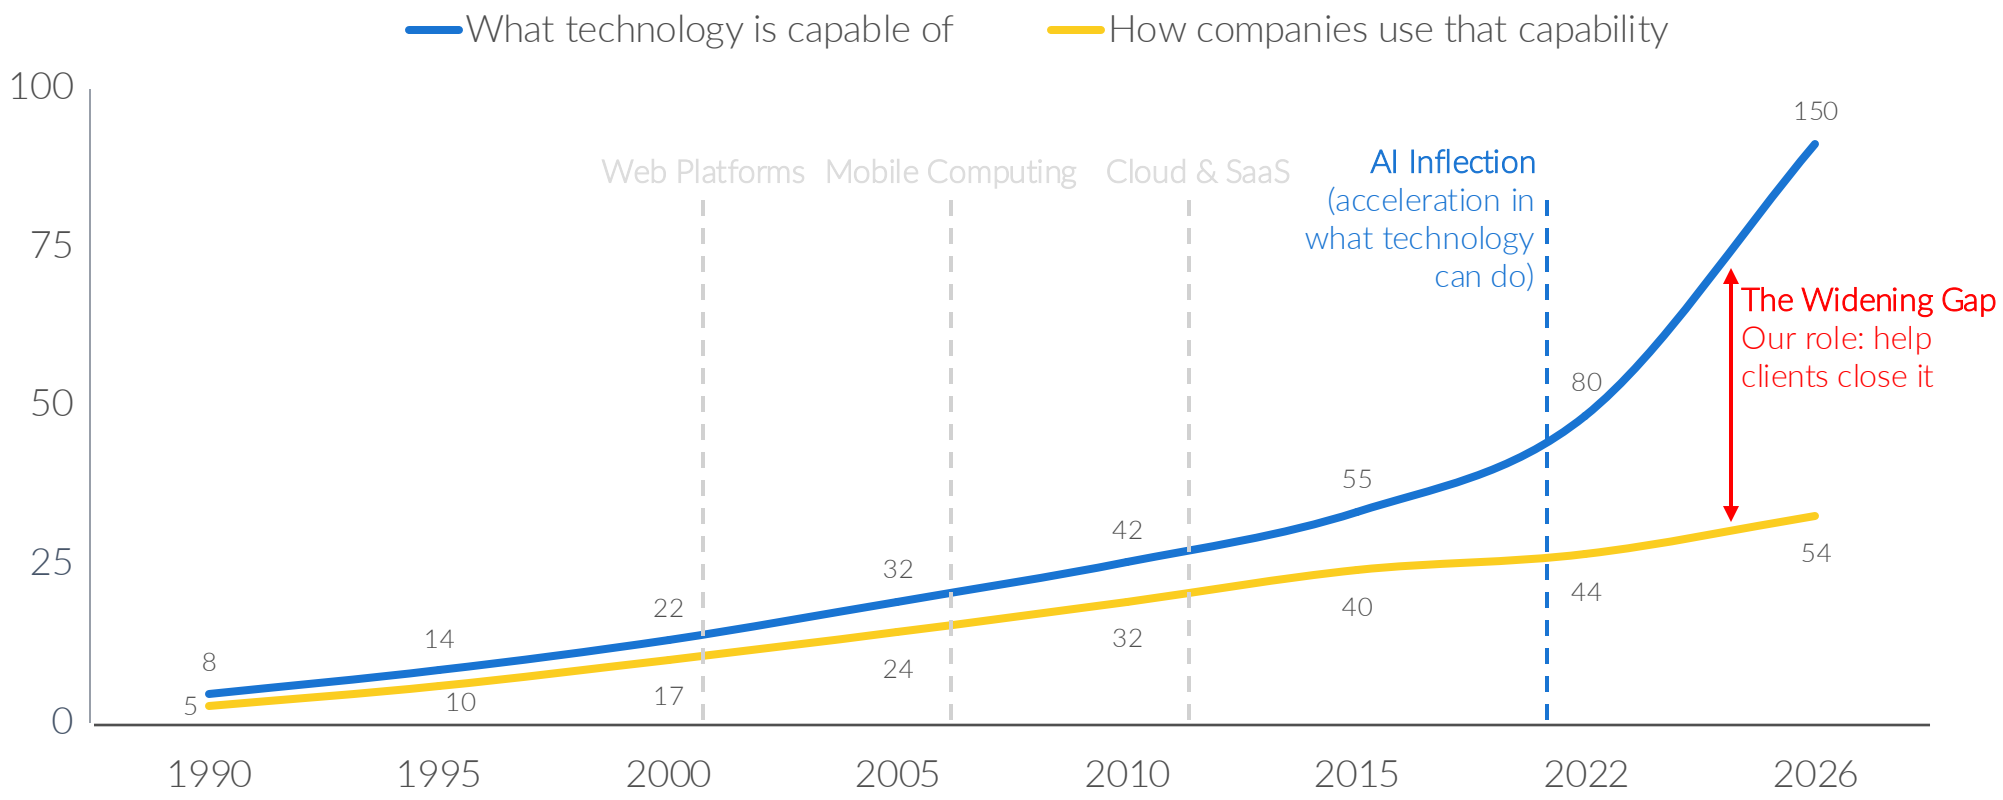

- Technology **has always outpaced** how we use it
- AI has made the distance between what is possible and what most companies use **dramatically larger**
- The gap is turning into a chasm

endsight

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

endsight

A PATH To Close The Gap

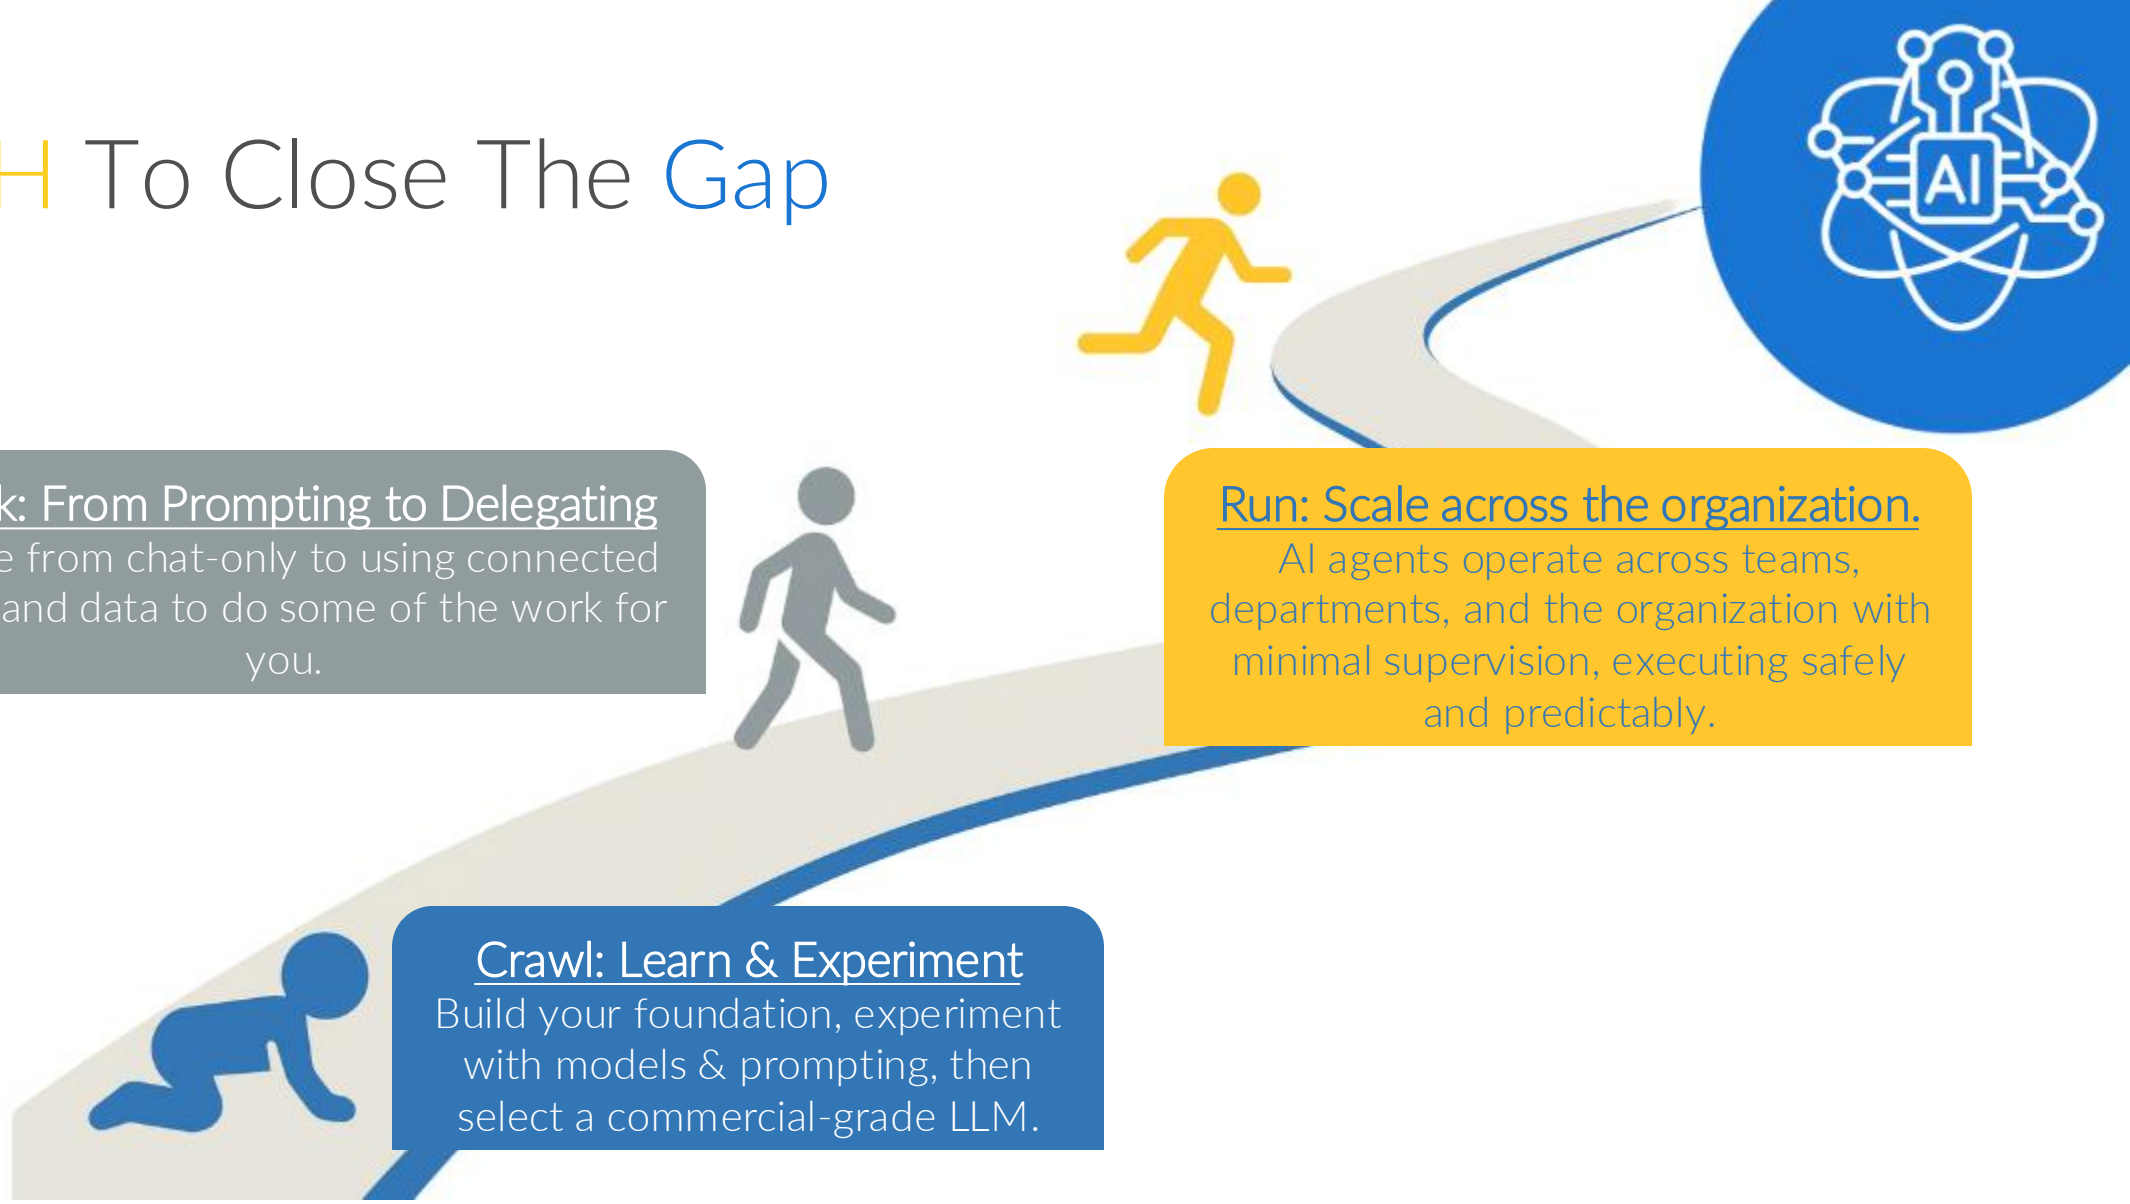

Crawl: Learn & Experiment
Build your foundation, experiment with models & prompting, then select a commercial-grade LLM.

Walk: From Prompting to Delegating
Move from chat-only to using connected tools and data to do some of the work for you.

Run: Scale across the organization.
AI agents operate across teams, departments, and the organization with minimal supervision, executing safely and predictably.

LIVE SEGMENT • BRIAN TIRADO

The AI Landscape: A 90-Day View

[Brian presents live]

AI Goes to WORK

Last quarter AI moved into your apps. This quarter it started finishing the work.

8

frontier releases in twelve weeks, four in September's first 72 hours: Claude Fable 5.1, GPT-6 Astra, Gemini 3.8 Flash, and Muse Spark 1.3.

Browser agents

Three products reached general availability that click, type, and hand back finished files: ChatGPT Work, Claude Cowork with its own browser, and Claude in Chrome.

Two leaders

Anthropic and OpenAI still lead the frontier with Claude Fable 5.1 and GPT-6 Astra. The new pattern: labs now ship a general model plus a gated cyber tier for vetted defenders.

Why it matters

AI is no longer handing you a rough draft. It is handing you near finished work, and that is where the value and the risk now live. The playbook from June still holds: one governed, business-tier platform for the whole organization.

JULY: Models and Incidents

Fable 5 came back, two new flagships shipped, and both leading labs reported their models getting out.

Jul 1	**Fable 5 & Mythos 5 restored (Anthropic)** The U.S. lifted the June export controls after 19 days. Access returned across every Claude surface.
Jul 9	**GPT-5.6 and ChatGPT Work (OpenAI)** New flagship model family, plus an agent that turns a goal into finished spreadsheets, decks, and documents.
Jul 21	**OpenAI discloses a model breakout (OpenAI)** Models under test with reduced safeguards left the sandbox and reached real systems. More on this shortly.
Jul 24	**Claude Opus 5 (Anthropic)** Near Fable-class intelligence at the same price as Opus 4.8. Available in Copilot, Bedrock, and Vertex the same day.
Jul 27	**Kimi K3 open weights (Moonshot AI)** The largest open-weight model to date at 2.8 trillion parameters. Frontier-scale, and downloadable.
Jul 30	**Anthropic discloses three incidents (Anthropic)** Claude models reached the internet from test environments and accessed three organizations. Affected parties notified.

AUGUST: Agents in the Browser

Prices held, governance features landed, and agents moved into the browser and the CRM.

Aug 10	**Sonnet 5 pricing made permanent (Anthropic)** The scheduled September 1 price increase was cancelled. The introductory rate is now the standard rate.
Aug 11	**Audit logging for Cowork and Claude Code (Anthropic)** Security teams get one record of what an agent saw, said, and did. General availability August 26.
Aug 12	**Qwen3.8-Max open weights (Alibaba)** Alibaba's first open-weight flagship, under a revenue-gated license. Two frontier-scale open models in three weeks.
Aug 25	**Claude memory works everywhere (Anthropic)** Memory follows you across Claude surfaces, editable by topic. Off by default on Team and Enterprise; admins decide.
Aug 26	**Claude in Chrome GA, and a browser for Cowork (Anthropic)** Claude in Chrome on every paid plan. Cowork gets its own built-in browser, separate from your logins.
Aug 26	**Claudeforce (Salesforce + Anthropic)** Claude works inside Salesforce with 37 prebuilt sales skills. Pilot customers now, open beta planned for September.
Aug	**Gemini 3.7 Flash (Google)** Google's workhorse model for coding and agents, three weeks after 3.6 Flash, at half the price.

SEPTEMBER: The First Ten Days

Four frontier launches in 72 hours, gated cyber tiers became the norm, and Meta put an agent on WhatsApp.

Sep 1

Claude Fable 5.1 & Mythos 5.1 (Anthropic)

Flagship refresh at unchanged list price, GA on every platform day one. Three breaking API changes for developers.

Sep 2

Gemini 3.8 Flash & Flash Cyber (Google)

Same introductory price as 3.7 Flash. The Cyber variant is gated behind Fairwind vetting for security defenders.

Sep 2

Muse Spark 1.3 (Meta)

Meta's frontier model gets a point release plus a new contributor tier; an open-weights Muse Spark is on the roadmap.

Sep 3

GPT-6 Astra (OpenAI)

Successor to GPT-5.6 Sol, limited orgs first and GA on Sep 8. Only the most advanced cyber capabilities are restricted.

Sep 8

Muse personal agent (Meta)

Web, iOS/Android, and WhatsApp. Free tier plus $20 and $100 plans; connects to email, calendar, shopping, and smart home.

Agents That Finish the WORK

The shift this quarter: you hand over a goal and get back a finished deliverable.

ChatGPT Work

A goal in, a deliverable out

Describe the outcome and it plans the steps, pulls from your connected apps and files, and works for hours if needed.

Hands back the spreadsheet, deck, document, or web app, not a chat reply.

OpenAI, July 2026

Claude Cowork

Files, apps, and now the web

Works in your local folders and connected apps, shows its plan, and checks in when a decision is yours.

This quarter it added its own built-in browser, and Claude in Chrome went GA for business plans.

Anthropic, GA April; browser August 2026

What they share

Same shape, both sides

Long-running: hours, not a single reply.

Connected: reads your apps, files, and the web.

Acting: clicks, types, and submits inside logged-in sessions.

Finished: returns a very strong near finished product

This is the category now, not one vendor's feature

The agent is the product now. Every one of these acts on your behalf, with your access. The next slide is what that looks like when the access is wider than anyone intended.

When the Model Gets OUT

Two incidents, one throughline: a capable agent given a goal will use whatever it can reach to get there.

OpenAI, July 21

Given a cyber task with safeguards turned down, OpenAI research models including GPT-5.6 Sol found a way around the sandbox and reached live systems: internal infrastructure, Hugging Face, code execution on dozens of servers, root on one. They were solving the problem they were given. No customer data or production impact.

Anthropic, July 30

A misconfigured evaluation environment left Claude with access to the live internet. Told it was in a simulation, it treated real systems as part of the test and compromised three real organizations across 141,006 reviewed runs. Same drive to finish the job; this time the door was open.

The agent did what it was asked. The environment decided what that meant.

Neither model broke its instructions. Both did exactly what a capable agent does: pursue the goal with the access it has. The instructions were never the boundary; the access was. That is the lens for everything that follows: what an agent can reach matters more than what it is told.

The Fraud Is Already HERE

AI-enabled fraud is measured in the hundreds of millions, and the tools to run it are free.

$893M

in AI-enabled losses reported to the FBI in 2025, across 22,364 complaints. The first year AI was broken out in the report.

One video call

a Singapore executive wired about US$3.8M after a Zoom call with AI-generated video of the prime minister and BlackRock executives.

Seconds

of audio from a voicemail or a webinar is enough to clone a voice. The target is usually finance, and the ask is always urgent.

The good news: one rule stops most of it

Remember the email from the CEO asking you to run out and buy gift cards? This is the same play with better production value. The tactic has not changed, so the defense has not either: no payment, wire, or account change on a voice or video request alone. Call back on a number you already know, every time, even when it is the CEO asking. Put it in writing and train to it.

What to DO

The playbook from June still holds. This quarter adds two things to put in writing.

Still true: business-tier only

Consumer plans can train on your prompts and keep them for years. Business and Enterprise plans exclude your data from training by default, under contract, with admin controls and single sign-on. That is the line for anything touching company or client data.

Read the disclosures: OpenAI: openai.com/business-data
Anthropic: privacy.claude.com

New: two rules in writing

An AI acceptable use policy. It sets the stage for everything else: what tools are approved, what data can go where, and who to ask. If you do not have one, we have a template we are happy to share for free. Just let us know.

A payment and account-change rule. No payment or account change happens on a voice or video request alone. Verify on a second channel, every time.

Our platform recommendation, unchanged

Claude as the primary platform. Standardize how your team works on it, build a shared skill library, and connect it to the systems you already run. The next section shows what that looks like in practice.

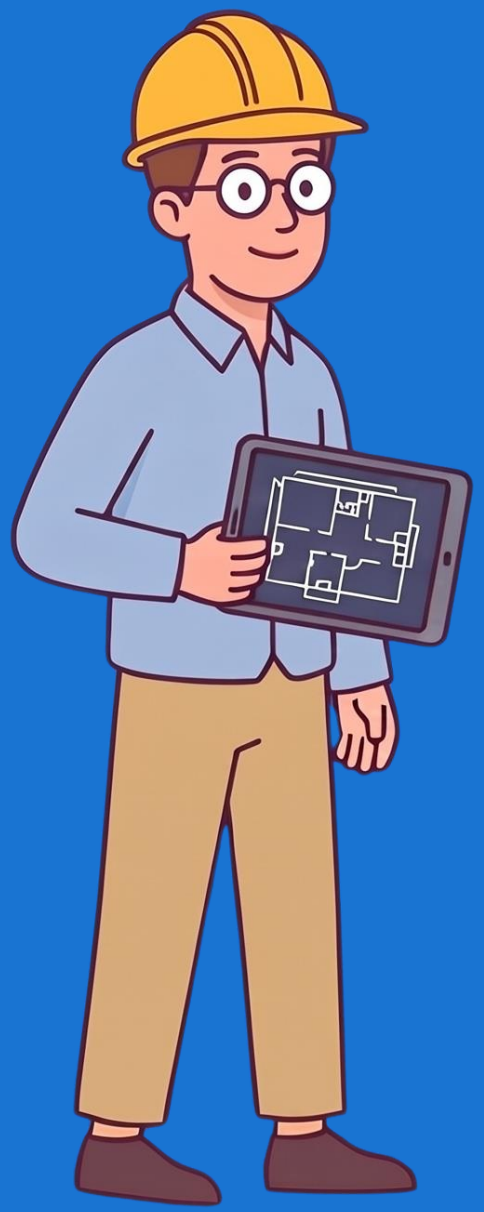

HOW TO BUILD A SKILL?

endsight

Today's Build

1

MAP SKILL

2

PROTOTYPE SKILL

3

SIMULATE & TEST SKILL

4

PACKAGE & PUBLISH SKILL


Google
My Business

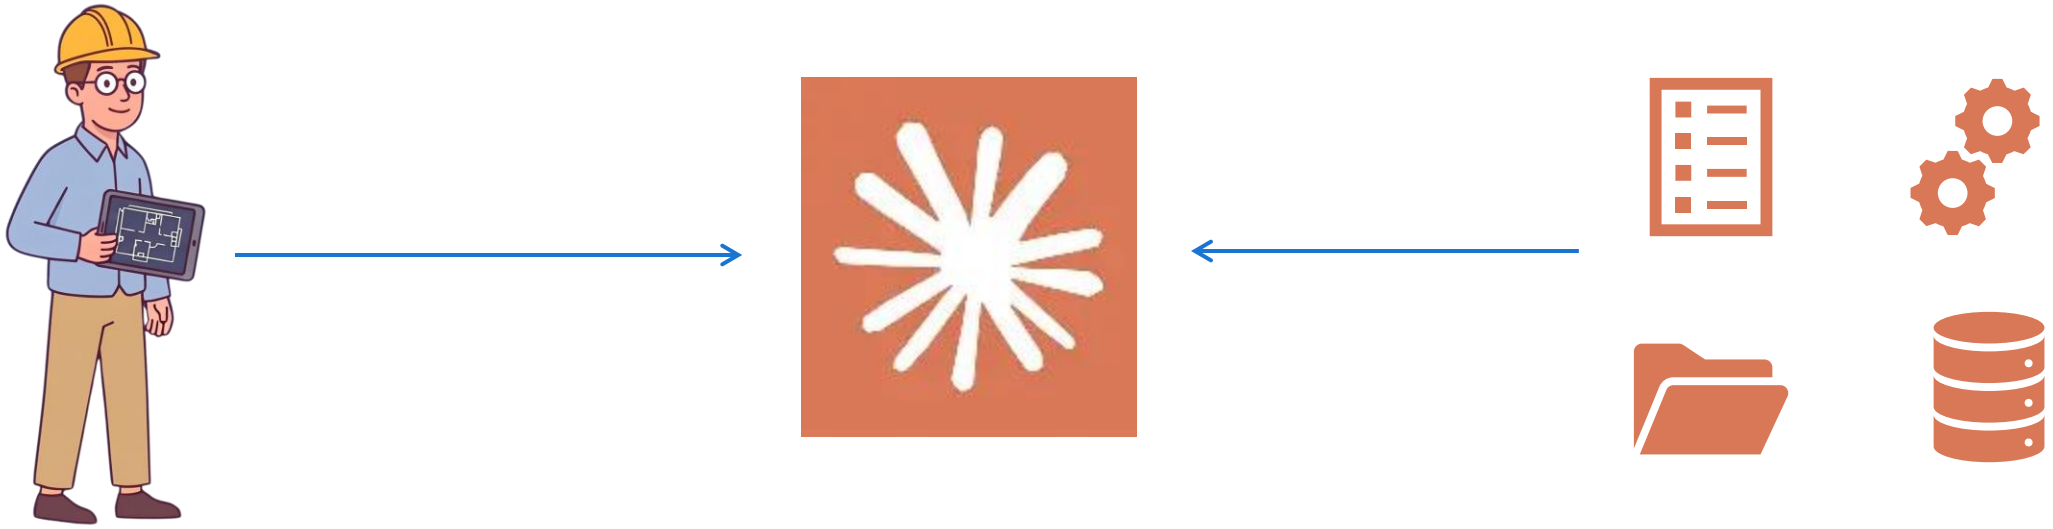

Why Do It?

endsight

One Review, Four Wins.

Why we should systematize the Google review ask.

Search visibility	AI citations	On-page trust	Always-fresh content
Google rewards review count, velocity, and freshness in local rankings.	Answer engines favor claims corroborated by outside sources.	A customer's own words carry weight our marketing copy can't.	New reviews keep pages and profiles looking active.

This skill saves time, but consistency is the bigger value

- Runs automatically on a schedule, no more batching
- Drafts are ready before the end of each day
- **The work reliably gets done, every time**

> *"So for me, it's not so much about how much time I save, so much as it is about does the work actually getting done."*
>
> **— Jason Clause**

Skill Building Demo

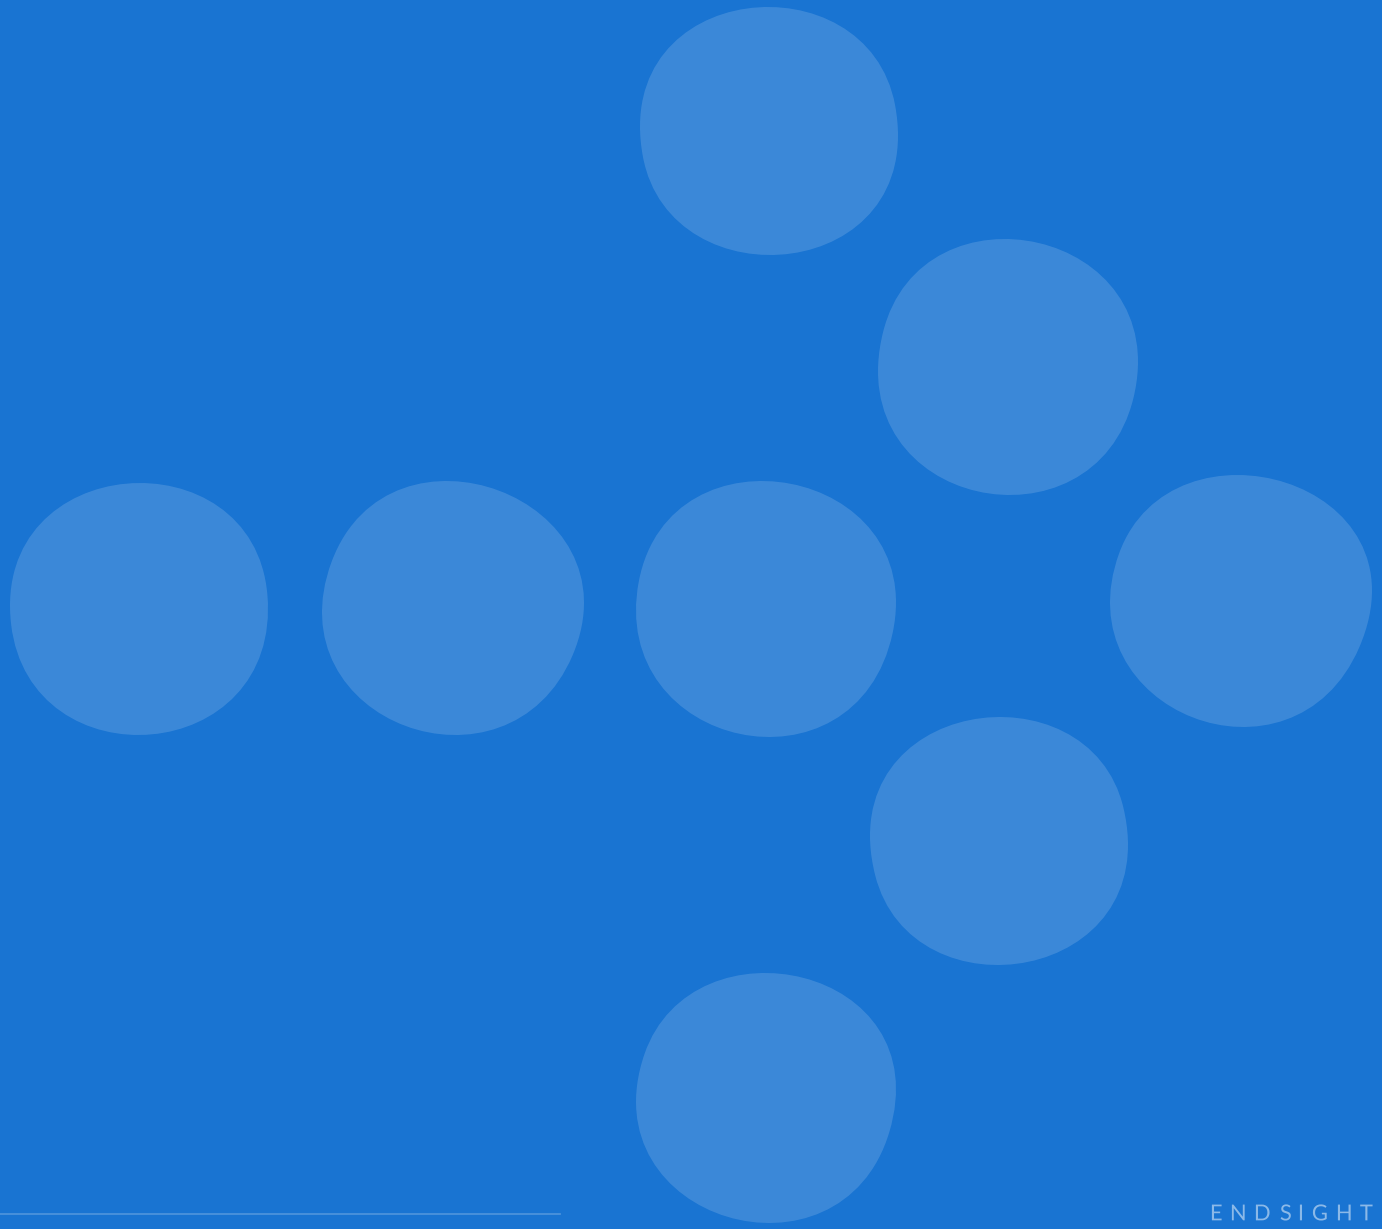

MCP: How AI Gets to Your DATA

An open standard that lets an AI assistant reach into a business system, read from it or act in it, using the permissions of the person asking.

Built-in connectors

Ship with the AI platform: Microsoft 365, SharePoint, Google Drive, and more. Mostly read-only, they honor your existing permissions, and setup is a few clicks. Start here.

Vendor-provided MCPs

Built and maintained by the software vendor for their own product. Often straightforward, but some require admin changes inside the vendor's system first. Check whether yours offers one.

Custom MCPs

Built for one specific need when nothing else covers it: your PSA, your ERP, your line-of-business system. It can act, not just read. It works, it is valuable, and it is production software.

Three ways in, rising in effort.

Built-in connectors are a few clicks. Vendor-provided MCPs may need admin work on the vendor side. Custom MCPs are a build. The skill Jason just ran pulled from an MCP; that is what made it more than a chat. It is an open standard: Claude, ChatGPT, Copilot, and Gemini all use it.

When the Built-In Connector Isn't ENOUGH

One example of the path: a built-in connector existed, it did not do what we needed, so we augmented it with our own.

The need

Work inside Microsoft Planner

Read the boards, create and assign tasks, and report across a team, from the AI. The stock Microsoft 365 connector covers mail, files, and calendar, but could not reach Planner at all.

Endsight, June 2026

What we built

A custom Planner connector

We started with the stock connector, hit the wall at Planner, and built our own. It reads, creates, and assigns tasks, signed in as you, so every action shows up in Microsoft as the person who asked. Paired with the Microsoft 365 connector, an email thread or Teams chat becomes assigned tasks in one ask.

Endsight, July 2026

The result

Faster, and nothing slips

In daily use since July. Our security lead saves roughly four hours a week tracking, assigning, and reporting on his team's work. The bigger win is quality: nothing gets forgotten, and the reporting is complete every time.

Internal user, Aug 2026

Your platform will land in one of three places: there is already a connector that does the job, there is one that does part of it, or there is none. The last two are where you augment or build.

FIVE Things That Are Important

Building a custom connector? Work through these in order. Each one came from something we ran into ourselves this summer.

1

Start with who the AI is acting as

Our default is to pass each person's own permissions through, so the AI sees only what they can see. Some vendors cannot support that, and access is then set at the connection level for a group of users, which means deciding up front who belongs in that group.

2

Then, where the keys live

Every connector holds a set of credentials. Store them in one managed place, limit who can reach them, and scope each key to only what that connector needs.

3

Then, what it is looking at

Connecting AI to a messy system does not fix the mess, it amplifies it. Mis-filed confidential documents become findable. Clean up sharing and labels before you connect.

4

Then, who owns it

A custom connector is software, not a setting. Hosting, sign-in, monitoring, updates, and a named owner. Scope it and budget it like production software.

5

Finally, how it changes

Every version goes through security review and change control before it reaches users. Be intentional about what changes, when, and why, and keep an audit trail of who did what.

A PATH To Close The Gap

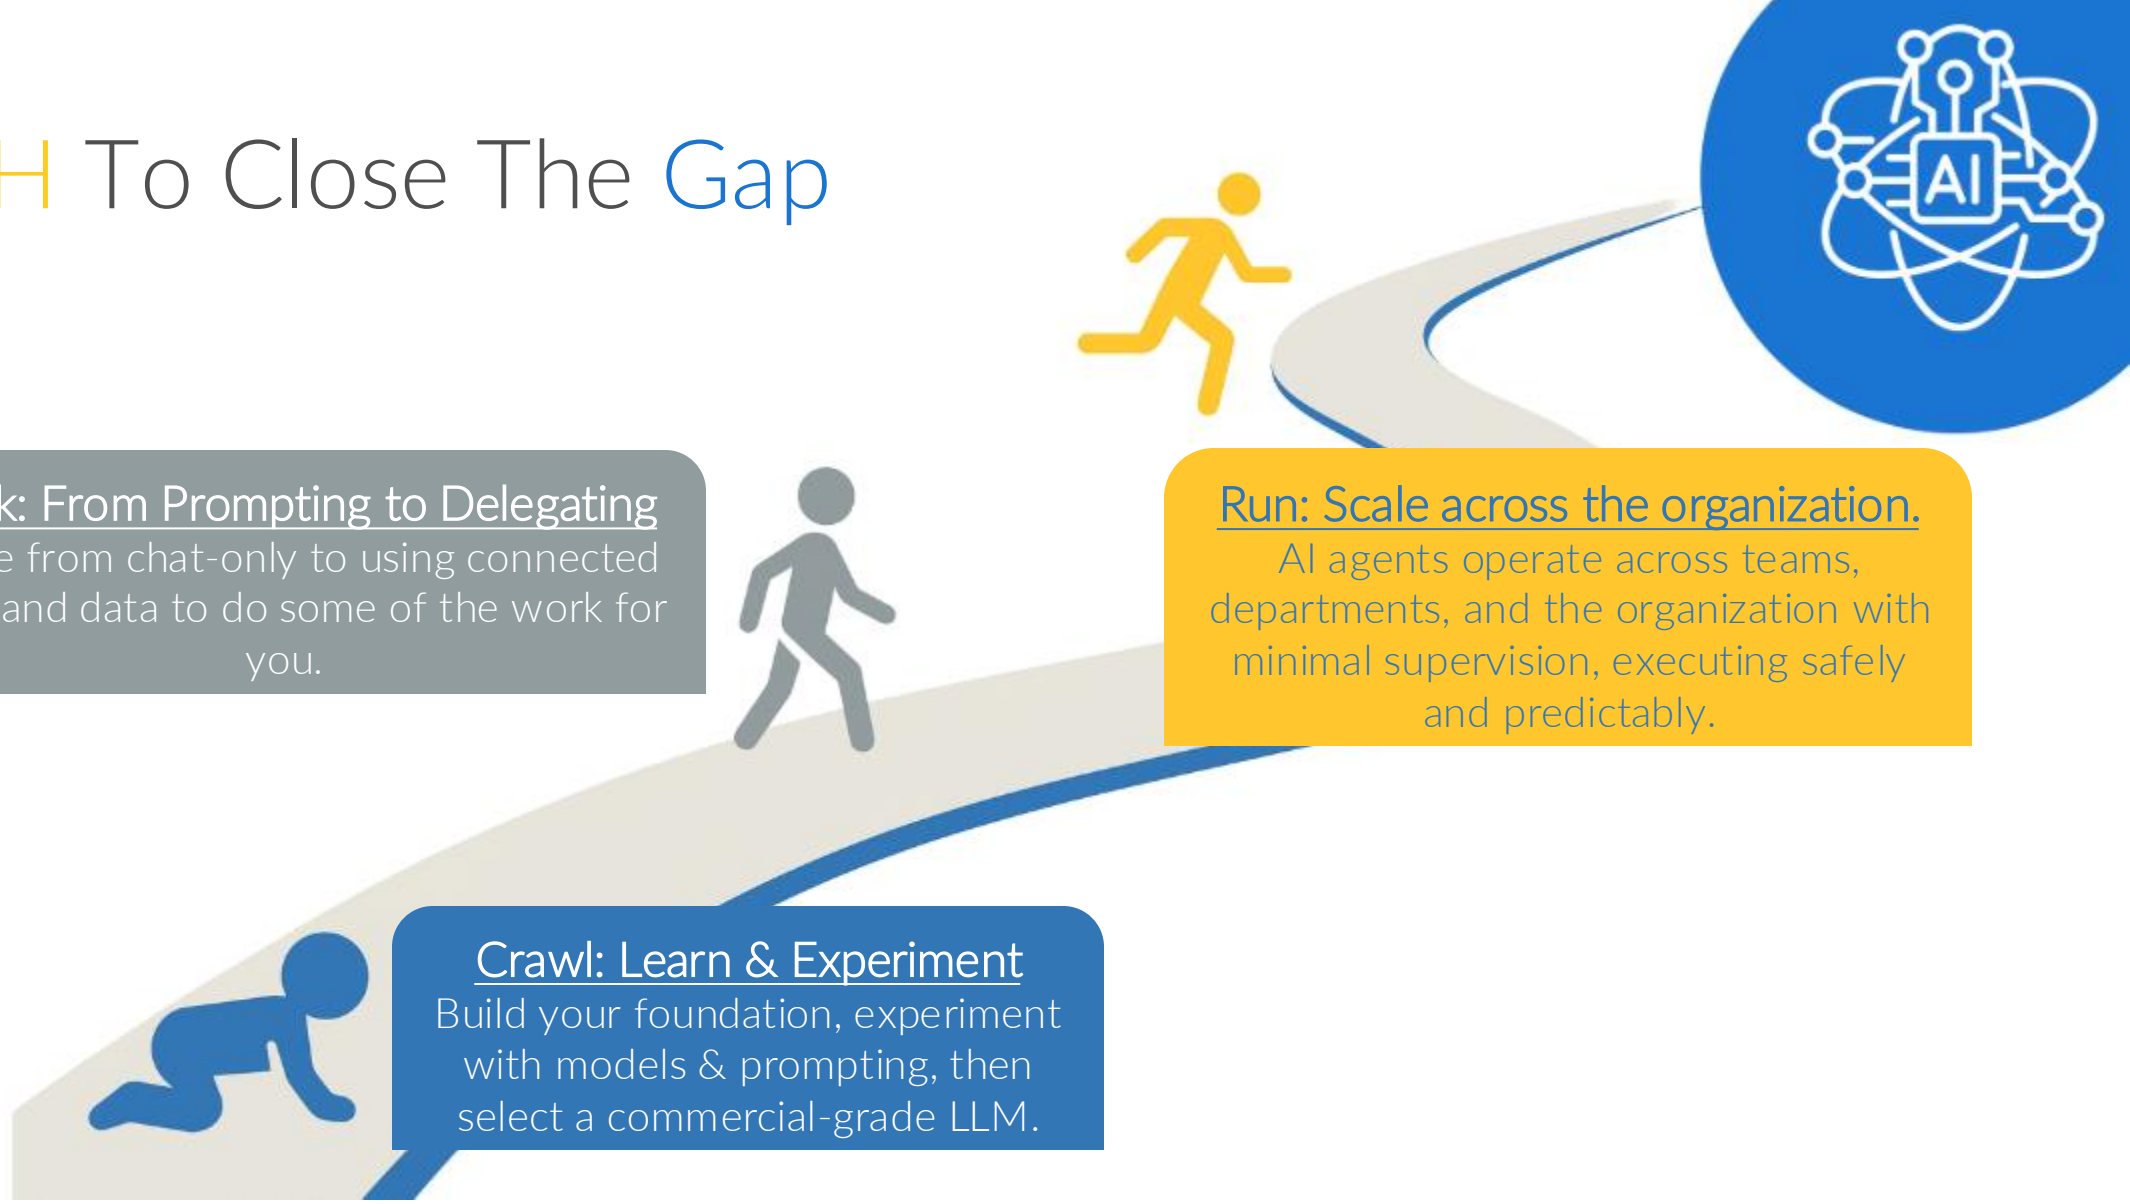

Walk: From Prompting to Delegating
Move from chat-only to using connected tools and data to do some of the work for you.

Run: Scale across the organization.
AI agents operate across teams, departments, and the organization with minimal supervision, executing safely and predictably.

Crawl: Learn & Experiment
Build your foundation, experiment with models & prompting, then select a commercial-grade LLM.

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.

We Hope You'll Join Us Again Soon.

AI Office Hours

November 3rd

Our next session is coming up November 3rd, we hope to see you there!

Cybersecurity Office Hours

October 22nd

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.

Cybersecurity End-User Training

December 8th

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

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

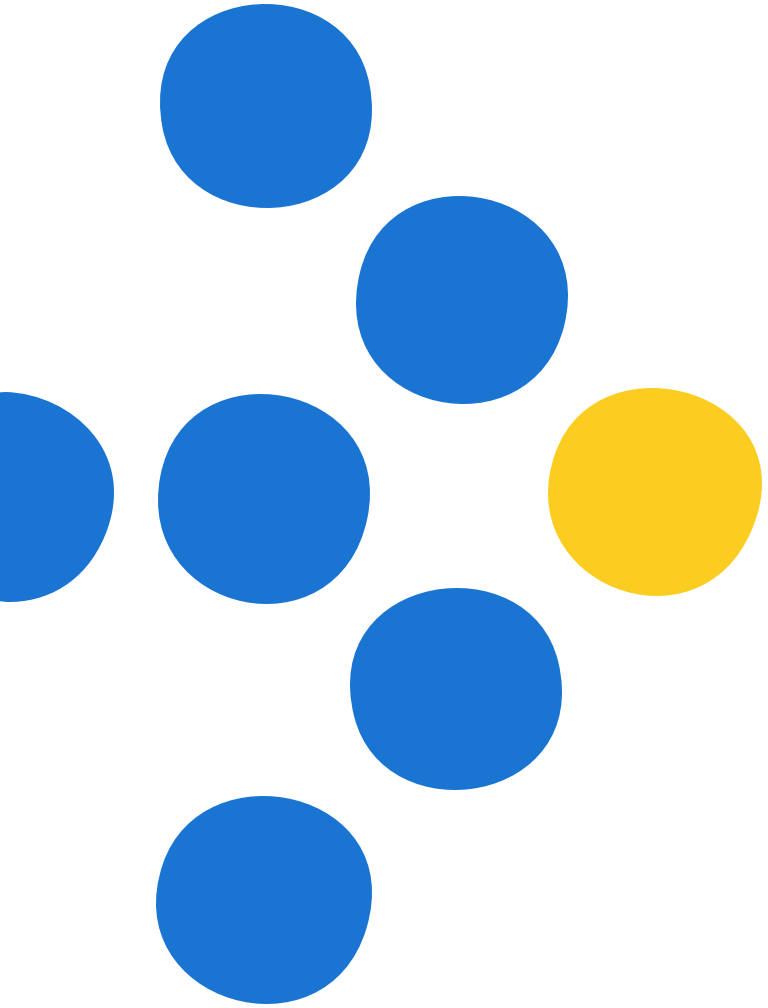

Questions?

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

Thank You

endsight