SAM BOURTON · GENEVA, SWITZERLAND

Under the Hood of Agentic AI

An executive primer — and everything you need to build your own. Take the deck, the setup guide, and the link with you.

WHO'S TALKING



A LITTLE ABOUT ME

```
name: Sam Bourton
based: Geneva, CH
building: AI · agents · second brains
```

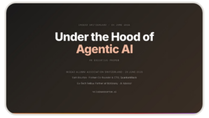
15 years building one of the world's leading AI practices — from a Formula 1 garage to a 1,500-person capability inside McKinsey. Today an AI founder and advisor, partnering with leaders to design and build with Progression Labs, our applied AI studio.

FIND ME

sambourton.ai →
my website

linkedin.com/in/sambourton →
LinkedIn

hello@sambourton.ai →
email



THESE SLIDES

The full deck (PDF)

Download every slide as a simple PDF to keep or share.

download PDF →



BUILD YOUR OWN

Setup guide

My Claude · Hermes · Obsidian setup, as markdown anyone or their agent can follow.

get the guide →

scan for the deck
insead.sambourton.ai

INSEAD SWITZERLAND · 24 JUNE 2026

Under the Hood of Agentic AI

AN EXECUTIVE PRIMER

INSEAD ALUMNI ASSOCIATION SWITZERLAND · 24 JUNE 2026

Sam Bourton · Former Co-founder & CTO, QuantumBlack

Ex-Tech Fellow Partner at McKinsey · AI Advisor

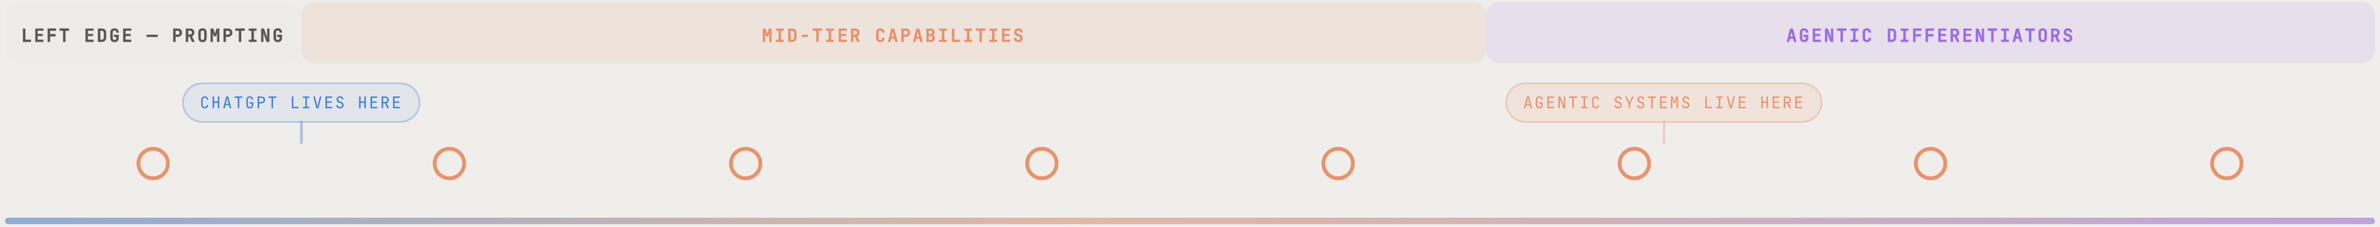
Progression Labs

hello@sambourton.ai

SLIDE 1 · THE AGENTIC SPECTRUM

Beyond the prompt

What many people think of as AI is the far left edge of a spectrum. Each step right adds capability and autonomy. Most companies live at step one. The leap is not a better prompt — it's a different architecture.



01
Prompting
Type a question, get an answer.

02
Context
Knowledge of you, your team, your business. Memory that compounds over time.

03
Connectors
Live links to email, calendar, drive, chat.

04
Tools
Acts on systems — Sharepoint, CRM, SAP, Salesforce.

05
Skills
Repeatable, reusable units of work.

06
Reasoning
Long-running, step-by-step decomposition.

07
Communications
Voice and chat as native interfaces.

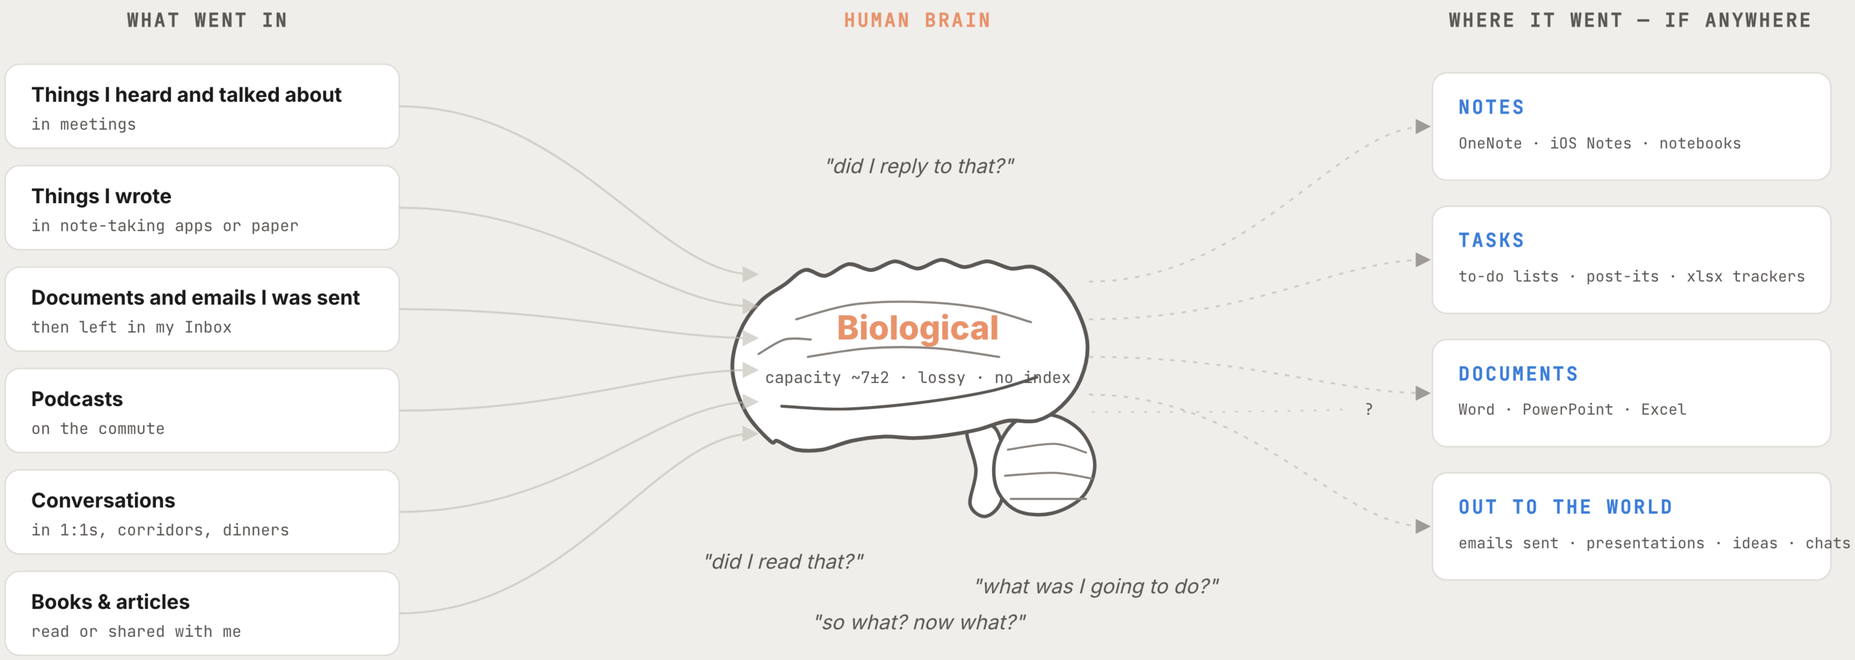
08
Outcomes
Give it a goal. It plans, acts, checks, loops until finished.

The strategic question is not "are you using AI?" Everyone is. It's **which steps along this spectrum are you operating at.**

SLIDE 2 · FIRST BRAIN · BIOLOGICAL

How I used to work

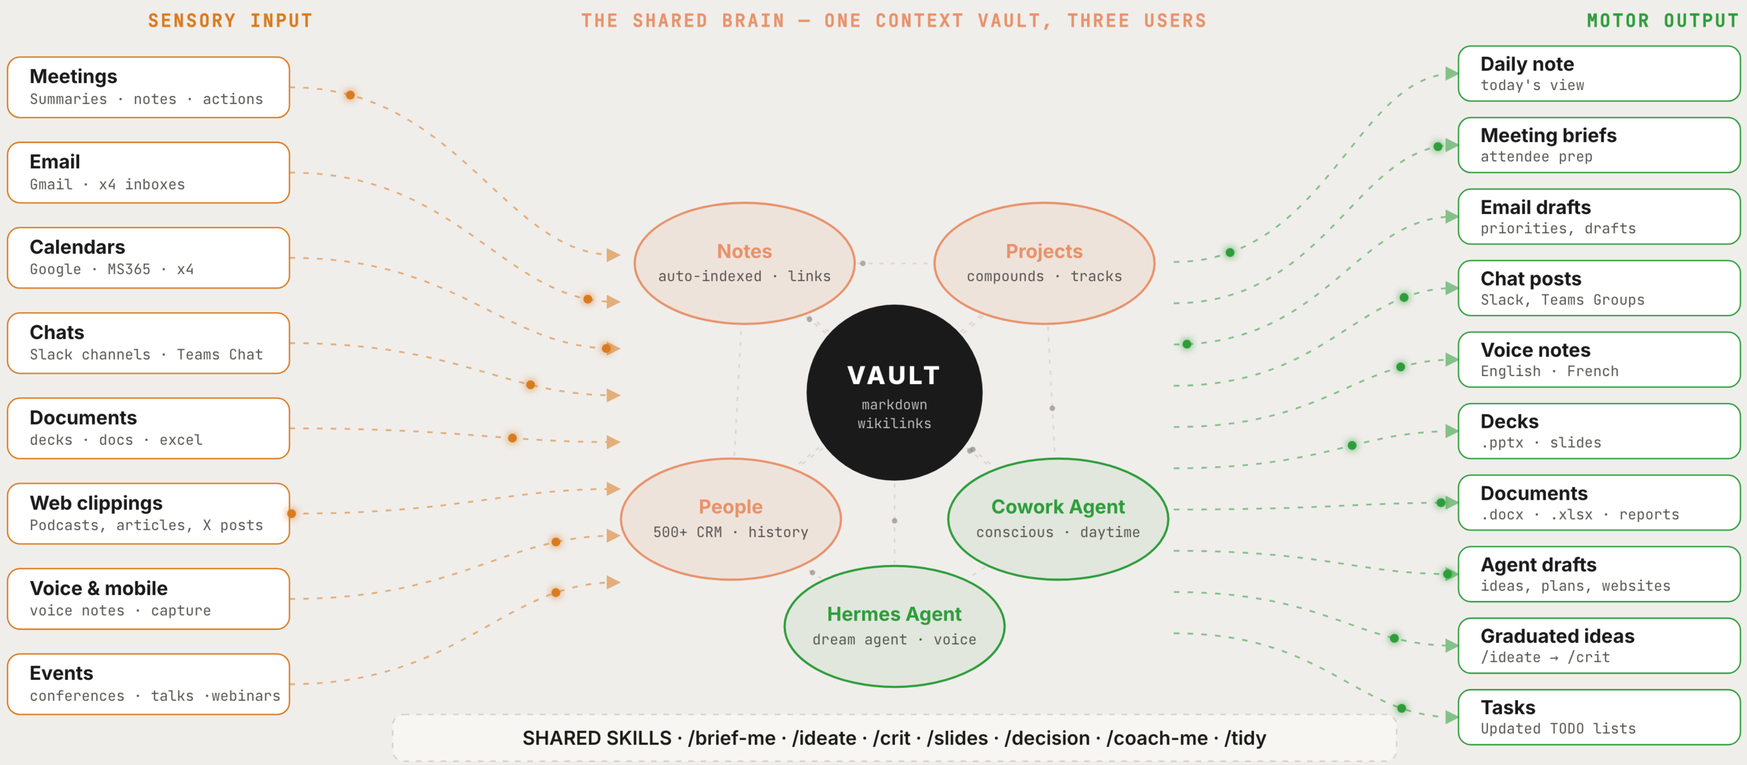
Analog in. Multiple systems and content stores. Scattered out and all manual. One biological brain in the middle trying to remember all of it — things drop off within 48 hours. The information was there; it just wasn't connected to anything.



What it cost: a meeting lived in my head for ~36 hours, then made it into one app or evaporated. Action items sat in three task tools. Every new deck or document was written from scratch.

The same inputs, wired together

Everything I see, say and touch flows into one shared vault. Three surfaces act on it — my human brain, Claude Cowork by day, Hermes by night and by voice. What comes out the other side is no longer information; it's decisions, auto-generated drafts and ideas.



Don't start by building an agent. Build the structure they will live in; the tools they can use; the skills they have; and the tasks you want to flow out. Start with simple tasks - research, daily briefs, regular status reports, drafting memos, creating powerpoints from your plans.

▶️ DEMO 1 · LIVE — SWITCH TO OBSIDIAN

The vault of shared context

Stop the slides. This is the real thing, not a diagram — the same vault every surface reads and writes. Walk through it live:

01
The knowledge graph
Everything linked — the shape of the second brain.

02
Project folders
How work is organised and compounds over time.

03
Clippings
Articles, posts and reading captured into the vault.

04
Meeting notes & task list
Auto-captured notes flowing into a live to-do list.

05
Daily briefs
What the agents generated overnight.

06
Email gravity dashboard
What needs me, ranked by what actually matters.

↩ When done, return to the deck for **Slide 4 · From me to the enterprise**

SLIDE 4 · AGENTIC MATURITY PATH IN THE BUSINESS

From the personal to the enterprise

What I built for myself is the embryo of what's about to happen to your organisation. There is a maturity path. Where you are on it is a strategic & business discussion, not a technology choice for IT.

START HERE ↓

MATURITY →

01 Virtual assistant

Every person has their own AI alongside them.

WHERE MOST ENTERPRISES ARE TODAY.

02 Virtual workflows

Discrete workflows owned end-to-end by AI, humans on the margins. The unit of automation is the workflow, not the role.

WHERE REAL ROI FIRST SHOW UP.

03 Virtual teammates

AI as a peer in the team — included in standups, owns deliverables.

WHERE THE EARLY MOVERS ARE HEADING.

04 Virtual teams

Whole teams of agents working together, with humans setting goals and reviewing outputs.

WHERE THE MANAGEMENT PROBLEM BITES.

05 Virtual business units

Functions like Finance or Procurement delivered as cross-functional agentic workflows — not org-chart units.

WHERE THE ORG CHART BREAKS.

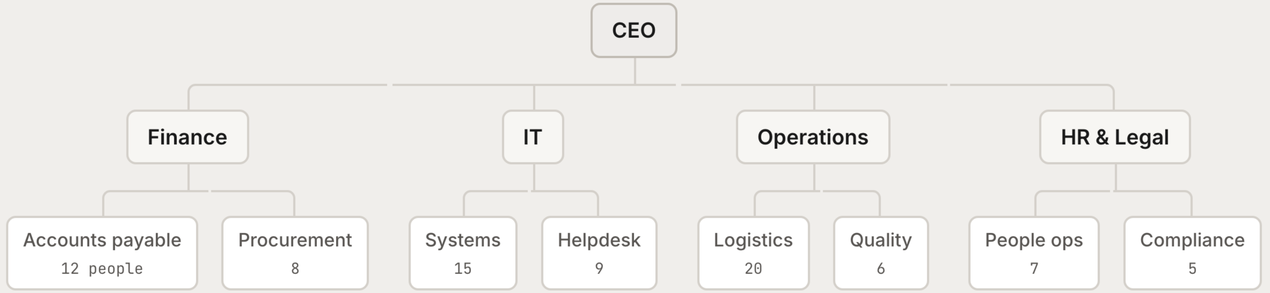
Examples I know first-hand. Progression Labs runs an agentic virtual CFO — not a single role replacement, but a cross-functional workflow spanning bookkeeping, reporting, forecasting and decision-support, with one human orchestrator on top. It does the work a small finance team used to do. **PhysicsX have a Strategy Assistant** — that runs 24/7 gathering competitor intel, job postings, news and announcements that may affect them. **Rezonant have a virtual Product Assistant** — synthesiszes all product data, research, insights, helps design features and products

We are transitioning from an era of AI experimentation to one of structural utility.

ORG CHART · PEOPLE IN BOXES

Two hundred years of managing people

Management theory was built for humans managing humans. The unit of management is headcount inside a function — managers who manage managers, in boxes and lines.

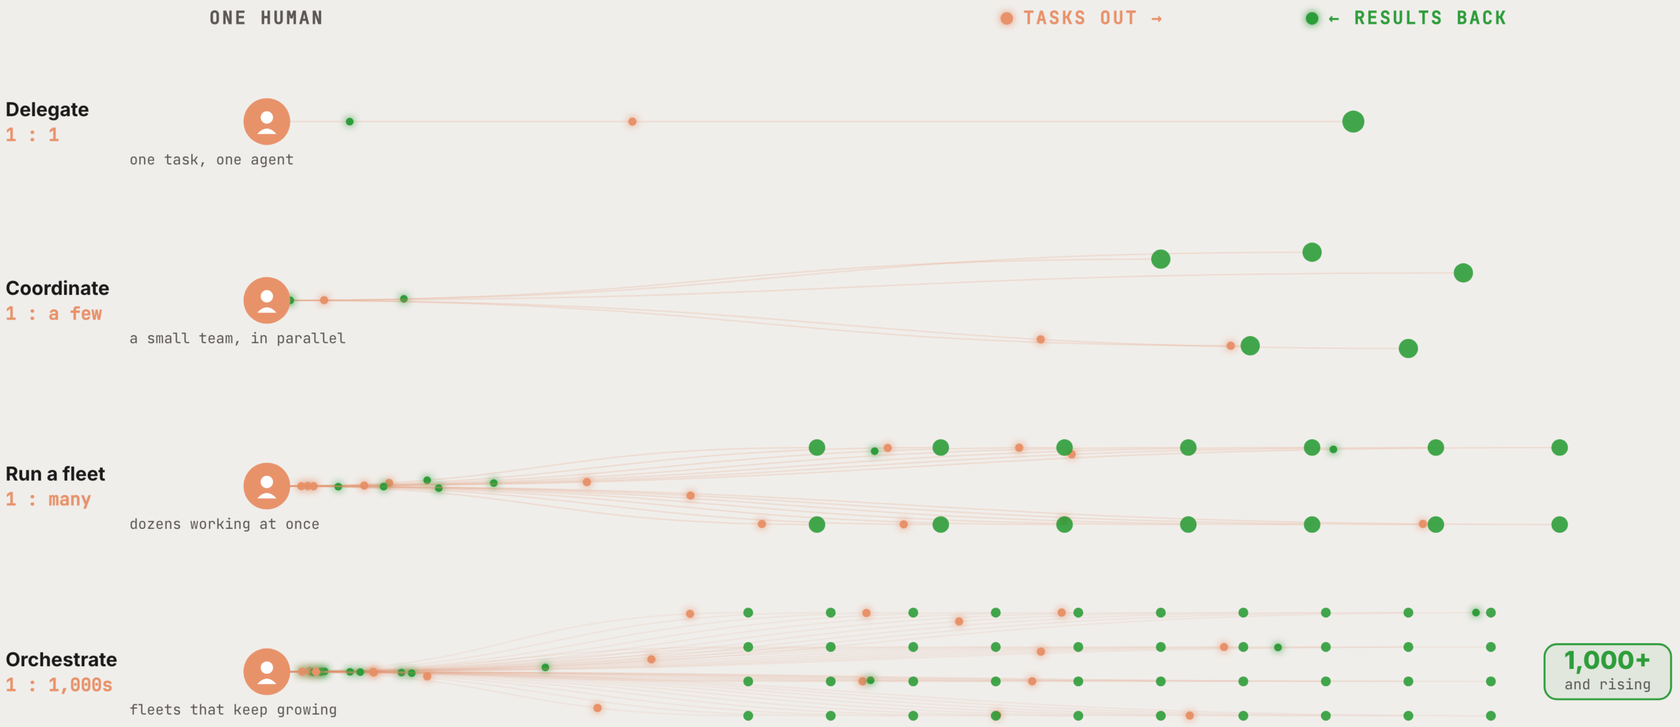


The unit of capacity is the role. You grow output by hiring more people into more boxes. For two centuries, that was the only way to scale. **The traditional org chart doesn't tell the whole story anymore.**

ORCHESTRATION · ONE HUMAN, MANY AGENTS

Now you orchestrate agents

The new unit of management isn't headcount — it's how many agents one person can direct. It starts one-to-one and doesn't stop: some operators already run hundreds, even thousands.



AI doesn't automate jobs — it automates tasks, and humans become orchestrators. The skill that compounds is no longer managing people. It's directing fleets of agents toward an outcome. Of every 100 jobs at the top AI companies, 14 exist to teach a model how your business actually works.

THREE MODELS OF HUMAN & AGENT COMMUNICATION

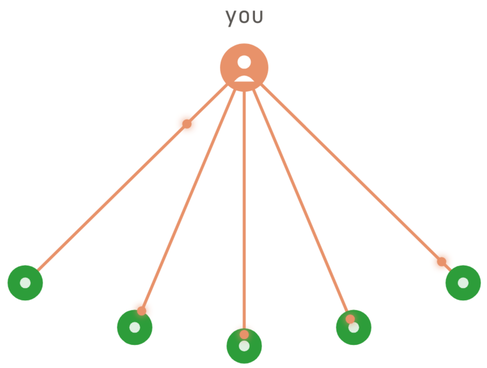
When agents start talking

As everyone gets their own agents, the real question becomes who talks to whom. Three models — each more powerful, and more tangled, than the last. Start from personal agents and progress only as your Trust, Guardrails, Risk, Security, ROI matures.

START HERE ↓

MATURITY →

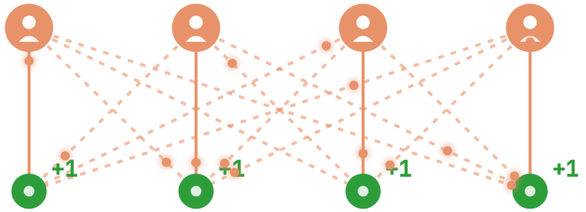
01 · TODAY



You and your agents

You talk to your own agents. One human directing their own fleet — the new model we just saw. AI never sends messages on your behelf.

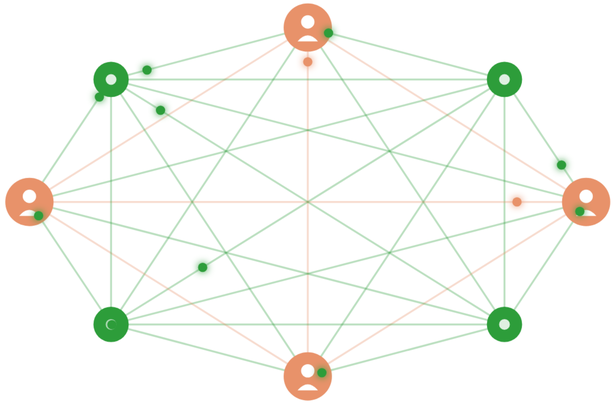
02 · PLUS ONES



Everyone gets a +1

Each colleague has a virtual **+1** that others can talk to — through the same chat and collaboration tools people already use. (The "Plus One" model.)

03 · AGENTS ↔ AGENTS



Agents talk to agents

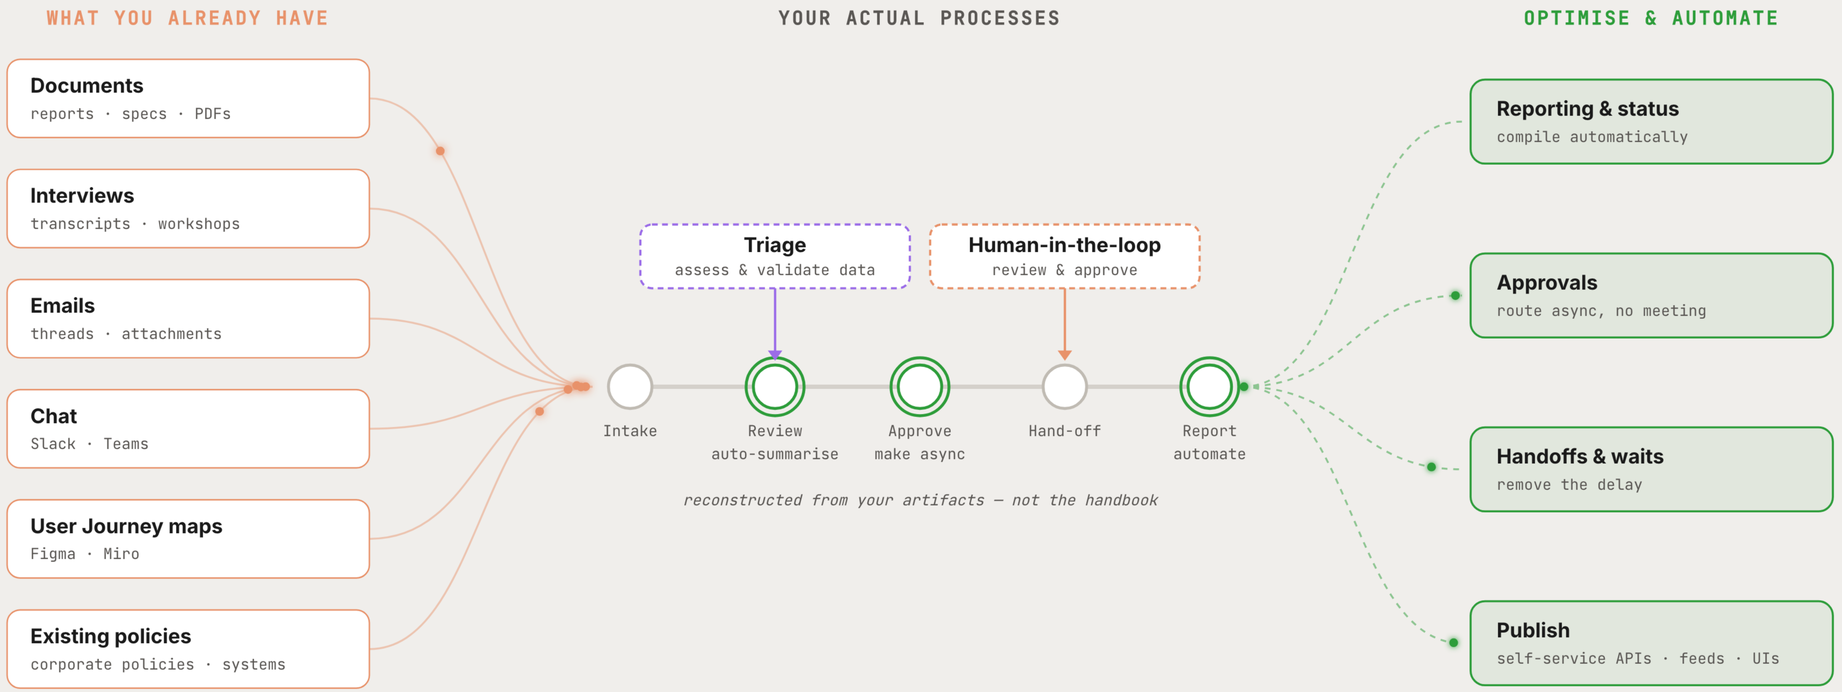
Your agents talk to other people's agents, then report back to humans. Hugely powerful — and complex very fast.

Your collaboration tools become the substrate for agent collaboration. The new governance question isn't "can we use AI?" — it's *who is allowed to talk to whom and about what.* Incrementally mature your use of agents "as fast as (safely) possible, but no faster.""

FROM YOUR ARTIFACTS TO A LIVE PROCESS MAP

Map what the org actually does

Start from the artifacts you already have — documents, interviews, emails, old process maps, policies, or map user journeys with Miro or Figma. Use AI to reconstruct the processes people really run, then flags where to optimise and automate, and stage incremental improvements to workflow steps.



Visibility before automation. Most organisations don't have a technology problem — they have an information-flow problem. Map the real work first; automate a broken process and you just break it faster. Use AI to auto-detect data validation issues at the point of intake, and suggest remediations. Route decisions to humans for final judgement.

2030 · BASE CASE

In 2030, every executive in this room will be running a fleet.

The base case for what working in a business looks like when agentic systems are normalised. Grounded in what credible voices and current evidence are pointing at — not utopia.

MEETINGS

Pre-briefed, transcribed, filed

Agents brief both sides, attend, log minutes and file follow-ups before anyone leaves the room. Roughly 30% of meetings cancelled — resolved async.

COMMS

Your inbox is an agent's inbox

Triage is fully delegated. You read the synthesis, not the threads. External comms is increasingly agent-to-agent.

DECISIONS

Brief, counter-brief, confidence

Every operational decision arrives with an agent-generated brief, counter-brief and confidence interval. Humans ratify by exception.

ORG CHART

Fewer boxes, more lanes

Flatter and barbelled. The unit of management is the workflow, not the headcount. The middle synthesis layer is dramatically thinner.

ROLES

Agent managers, eval owners

Workflow architect, policy engineer, prompt librarian are common titles. Most middle-management synthesis roles consolidated.

SERVICE BUSINESSES

Hybrid software companies

Consulting, law and life sciences sell judgment and eval at the top, agent-delivered work at the bottom. Fixed-fee replaces T&M.

88%

ENTERPRISE AI PILOTS THAT FAIL TODAY

MIT NANDA, 2025

171%

AVERAGE ROI FOR THE 12% THAT SUCCEED

McKinsey, 2026

14.5h

LONGEST AUTONOMOUS AGENTIC-INFERENCE RUN

Claude Opus 4.6, 2026

"You can outsource your thinking. You cannot outsource your understanding."

— Andrej Karpathy · Sequoia Ascent 2026

▶️ DEMO 2 · LIVE — ANSWER THE AUDIENCE WITH MY VOICE AGENT

Ask the room. Answer it together.

Take live questions from the audience — what they still want to know about AI — then answer them with me, using my voice agent over the shared vault and its agents.

01
Take questions
Ask the room what they still want to know about AI.

02
Feed the voice agent
Speak the questions straight into the voice agent, live.

03
It reads the vault
The agent searches my shared context and its agents for the answer.

04
Answer together
It answers out loud; I add the human half of the reply.

05
"What could I say?"
Ask it what, from my context, would land with this audience.

06
The point
Vault + agents + me — answering live, not from a script.

↩ When done, return to the deck for **Slide 7 · What you can do next**

DIAGNOSTIC FRAMEWORK · TAKE INTO Q&A

The questions you should be asking

Off the main path, but the framework to take into Q&A and back to your teams. Stop asking "should we do AI?" — everyone is. Start asking the four that actually matter.

QUESTION 01 · THE SPECTRUM

Where on the spectrum are we — and where do we need to be?

Map every workflow against the eight steps. Be honest about which ones are still at step one.

QUESTION 02 · THE PLUMBING

What's the deterministic plumbing around the model — and who owns it?

The model is the cheap part. The pre- and post-processing is where the engineering and the accountability live.

QUESTION 03 · END-TO-END OWNERSHIP

What workflow could an agent own end-to-end in six months — and what would have to be true?

Pick the workflow. Name the data, governance, and trust conditions. That's the real roadmap.

QUESTION 04 · SHAPER OR TAKER

Are we a Shaper or a Taker?

Are we building proprietary agentic ways of working — or only buying the tools everyone else buys?

*Then ask the same questions of **your vendors**. Then of **yourselves**.*

PRACTICAL TAKEAWAYS · THREE MOVES

What you can do next

Three moves to take from this room into your week. Each one is small enough to start before Friday and big enough to compound.

01

Use it personally

Get outside your comfort zone. Experiment. Learn by doing. You will not get the strategic instinct for this technology by reading about it — pick one workflow this week and try.

02

Get comfortable with security

Don't rush. Always be mindful, and slow down if you're unsure. Have practical conversations with your team and your IT function — work through the implications before you scale.

03

You can just ask Claude

A Claude Cowork + knowledgebase, set up in fifteen minutes, is enough to give a non-technical user serious leverage. *(My 87-year-old friend Lynne, the writer, did exactly that.)*

Get these slides and take a setup with you. Scan the code for the full deck — and from there, the link to my Claude · Hermes · Obsidian setup guide to seed your own second brain.

insead.sambourton.ai