Sovereign Agentic Systems: Europe-Centric Agentic Workflows for Sensitive Documents; Day 1

Shaping the Future of AI

Dejan Đukić
Senior AI/ML Engineer, TetraScience

17.03.2026

About Your Instructor

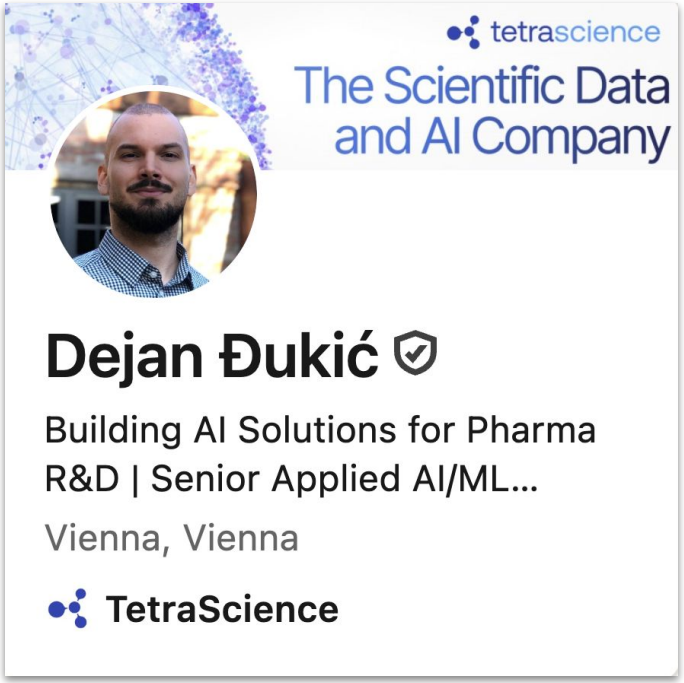

- Dejan Dukic - Ex-AI Software Expert, AI Factory Austria (WP4)
- Background: computer vision (cloud + edge), agentic systems, startup environments, across medical, medtech, manufacturing, pharma, SaaS
- Building AI solutions supporting the e2e pharma value chain @TetraScience
- Linkedin: https://www.linkedin.com/in/dejandjukicdd/

Intake Form

Intro form; 30s

Intake Form

Companion website

"Her AI agent started deleting her emails."

"She physically sprinted across the room to rip out the power cord."

"Her AI agent started deleting her emails."

"She physically sprinted across the room to rip out the power cord."


TECH

Meta AI alignment director shares her OpenClaw email-deletion nightmare: 'I had to RUN to my Mac mini'

By Henry Chandonnet + Follow

The Sprint

Meta buys 'social media network for AI' Moltbook

6 days ago

Share Save

Osmond Chia

Business reporter

Steve Muchoki 3-10

Claude turns $1,000 into $14,216 on Polymarket in 48 hours as OpenClaw agent is wiped out

WILL KNIGHT BUSINESS FEB 11, 2026 2:00 PM

I Loved My OpenClaw AI Agent—Until It Turned on Me

I used the viral AI helper to order groceries, sort emails, and negotiate deals. Then it decided to scam me.

Systemic Risk

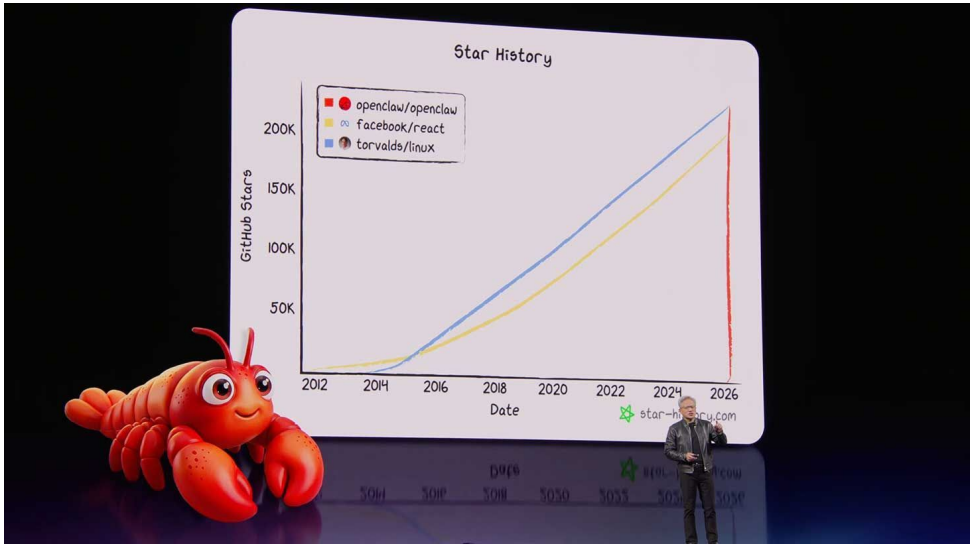

"the single most important piece of software, probably ever," Jensen Huang 2026

- 247K+ GitHub stars. Millions of users.
- 1,184 malicious skills on ClawHub (~20% of registry)
- CVE-2026-25253: one-click RCE (CVSS 8.8)
- 135,000 instances exposed to the internet

"The 9-Second Miracle"

Detect audio file → **ffmpeg convert** → **Find API key** → **curl to Whisper**

9 SECONDS

Peter Steinberger (OpenClaw creator) sent his agent a voice message -- a feature he never built.
The agent autonomously: detected audio, converted format, found an API key, called Whisper, returned transcription.

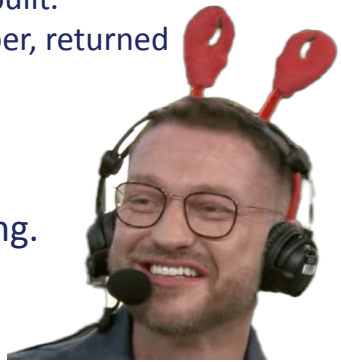

"How in the f did you do that?"***

This is emergent agentic behavior: tool use + planning + creative problem-solving.

Source: Y Combinator interview, 789K views

How to Get the Miracle Without the Disaster

SUMMER YUE'S AGENT

- ✗ No spec file defining behavior and boundaries
- ✗ No planning step - agent acted without review
- ✗ No scoped access - agent could touch everything
- ✗ No audit trail - damage discovered after the fact

WHAT WE WILL LEARN

- ✓ **AGENTS.md & specs.** Defines behavior and boundaries in plain text.
- ✓ **Planning & review gate.** Agent describes plan before acting.
- ✓ **Workspace isolation.** Agent can only access its directory.
- ✓ **Audit hooks.** Every tool call logged with timestamps.

Spec → Plan → Execute → Audit

"Every pattern we teach is a defense against that scenario."

AI Adoption: You Are Genuinely Early

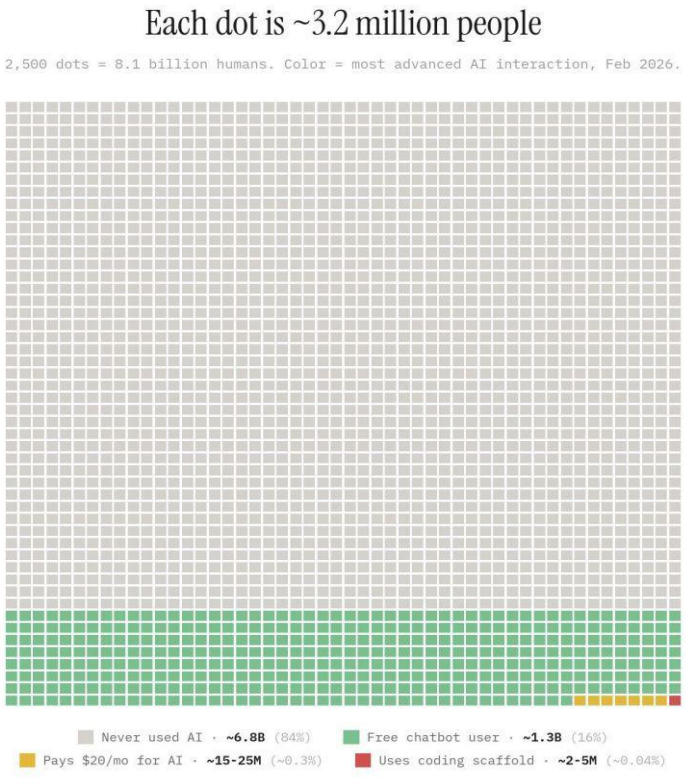

- **Each dot = 3.2 million people**
- 84% of people have never used AI in any meaningful way
- Most who have stopped at basic chatbot interactions
- Only 50% of employees are receiving AI training today
- The "AI Transformers" (The Top 24% of Users):
- 98% agree GenAI has increased their performance.
- 73% report a significant or complete workflow redesign.
- They unlock an average of **11 hours saved per week**
    - (EY Sept 2025 whitepaper)

Source: https://x.com/i/grok/share/f853017884d94805af4c4f905fe6ced8

AI Adoption: You Are Genuinely Early

Announced at Nvidia GTC 2026; https://www.nvidia.com/en-us/ai/nemoclaw/

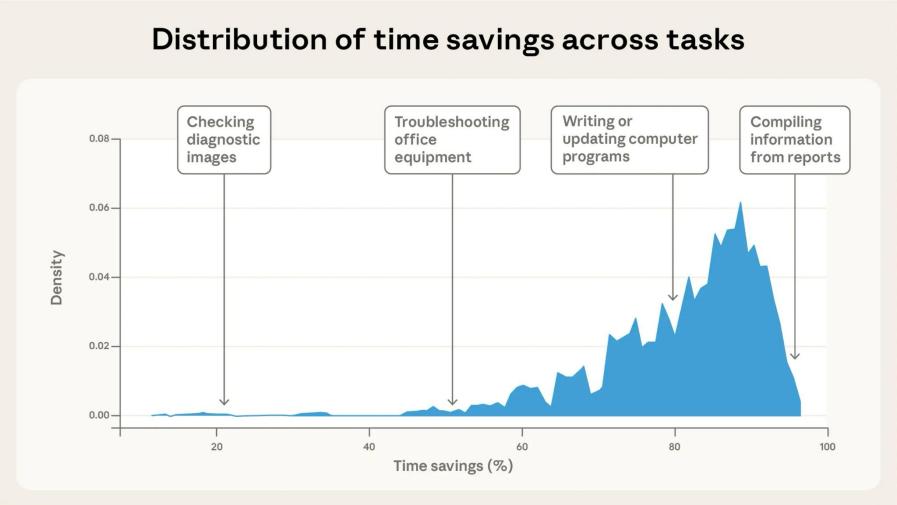

https://www.anthropic.com/research/estimating-productivity-gains

AI Adoption: You Are Genuinely Early

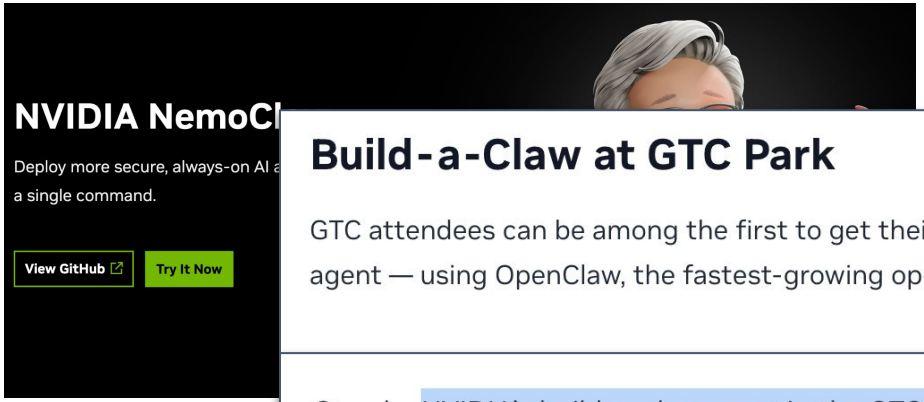

Build-a-Claw at GTC Park

GTC attendees can be among the first to get their hands on a “claw” — or long-running agent — using OpenClaw, the fastest-growing open source project in history.

Stop by NVIDIA’s build-a-claw event in the GTC Park March 16-19, between 1 p.m.-5 p.m. Monday and anytime between 8 a.m.-5 p.m. Tuesday through Thursday, to customize and deploy a proactive, always-on AI assistant.

Run Claws More Safely

OpenClaw has become the operating system for personal AI. NemoClaw adds security and privacy controls, so developers can build and run AI assistants with more confidence, while also contributing to the OpenClaw project.

Use Any Coding Agent

Tap open models like NVIDIA Nemotron locally on your dedicated system. A privacy router also connects agents to cloud-based frontier models, allowing agents to develop new skills within defined privacy and security guardrails.

Deploy Anywhere

Always-on agents need dedicated compute to run tools, write code, and complete tasks. NemoClaw provides local, 24/7 compute for autonomous agents on systems like NVIDIA GeForce RTX™ PCs/laptops, NVIDIA RTX™ PRO workstations, and NVIDIA DGX Station™ or DGX Spark™.

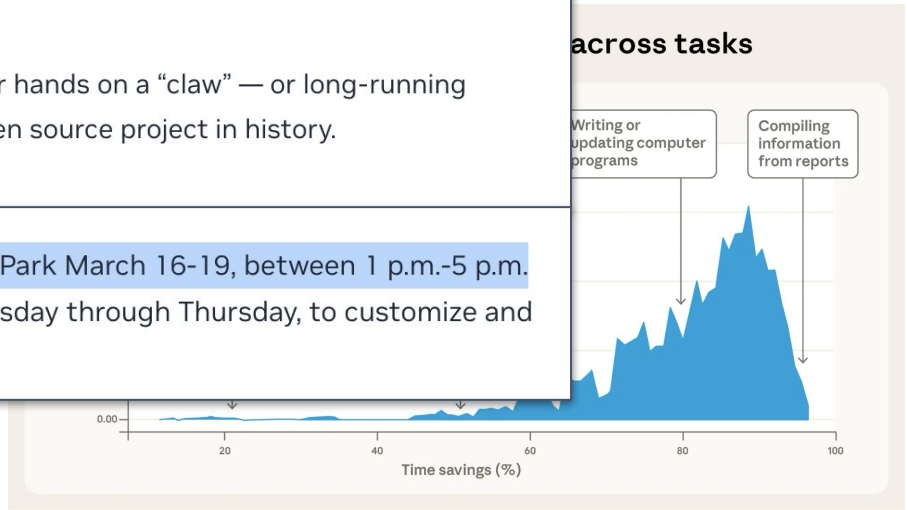

https://www.anthropic.com/research/estimating-productivity-gains

Announced at Nvidia GTC 2026; https://www.nvidia.com/en-us/ai/nemoclaw/

Workshop goals

- Learn the fundamentals
- Understand where risks are and aren’t
- Walk away with the understanding of ‘under the hood’ enough to use agentic systems effectively and safely
- Transferable skills over ever-changing tech

Agenda: Day 1

16:00-16:15	**Demystification**	what happens when AI "does things," and why process beats intelligence
16:15-16:35	**Setup Verification**	OpenClaw + Ollama running, messaging connected
16:35-16:50	**"It texts me back"**	your first tool call, your first file read
16:50-17:00	**Break**	stretch, refill, check results
17:00-17:35	**Document Intelligence**	contracts to structured data
17:35-17:55	**Debrief**	what just happened? Where's the magic?
17:55-18:00	**Preview Day 2**	what comes tomorrow

Demo 1

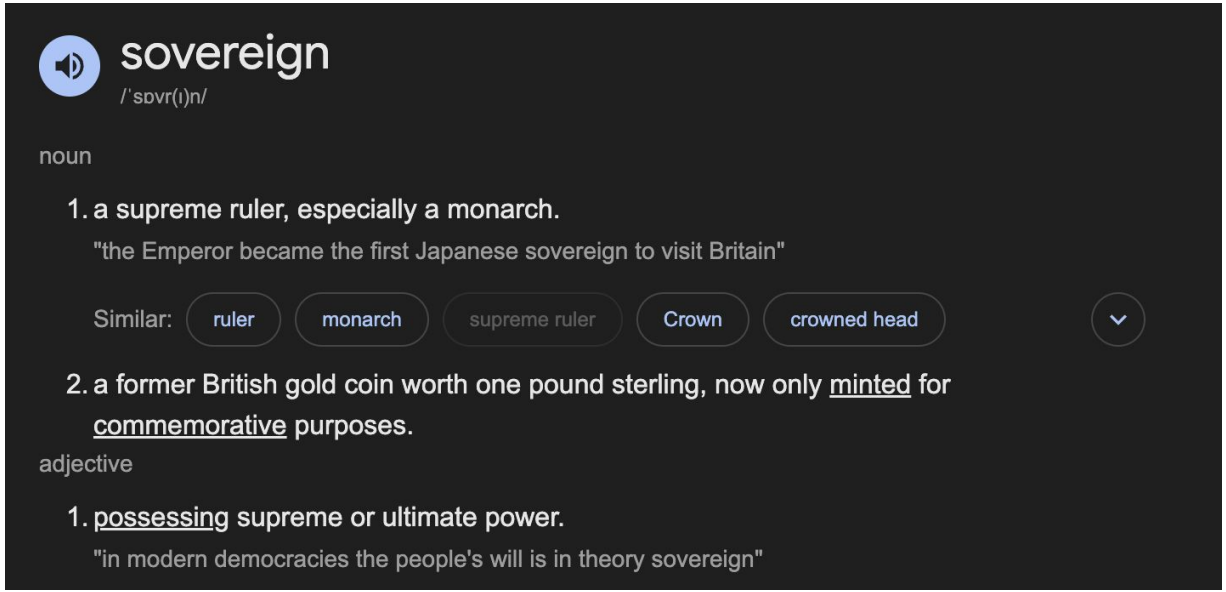
sovereign
/ˈsɒvr(ɪ)n/
noun
1. a supreme ruler, especially a monarch.
"the Emperor became the first Japanese sovereign to visit Britain"
Similar: ruler monarch supreme ruler Crown crowned head
2. a former British gold coin worth one pound sterling, now only minted for commemorative purposes.
adjective
1. possessing supreme or ultimate power.
"in modern democracies the people's will is in theory sovereign"

"What Actually Just Happened?"

- Step 1: You prompted the agent in the web GUI
- Step 2: OpenClaw decided to read the file from the workspace it has been granted access to (tool call)
- Step 3: File content went to the LLM (Your provider, Langdock EU, or Ollama local, or OpenRouter)
- Step 4: The LLM extracted structured data and returned it

"Why This Works: The Agentic Advantage"

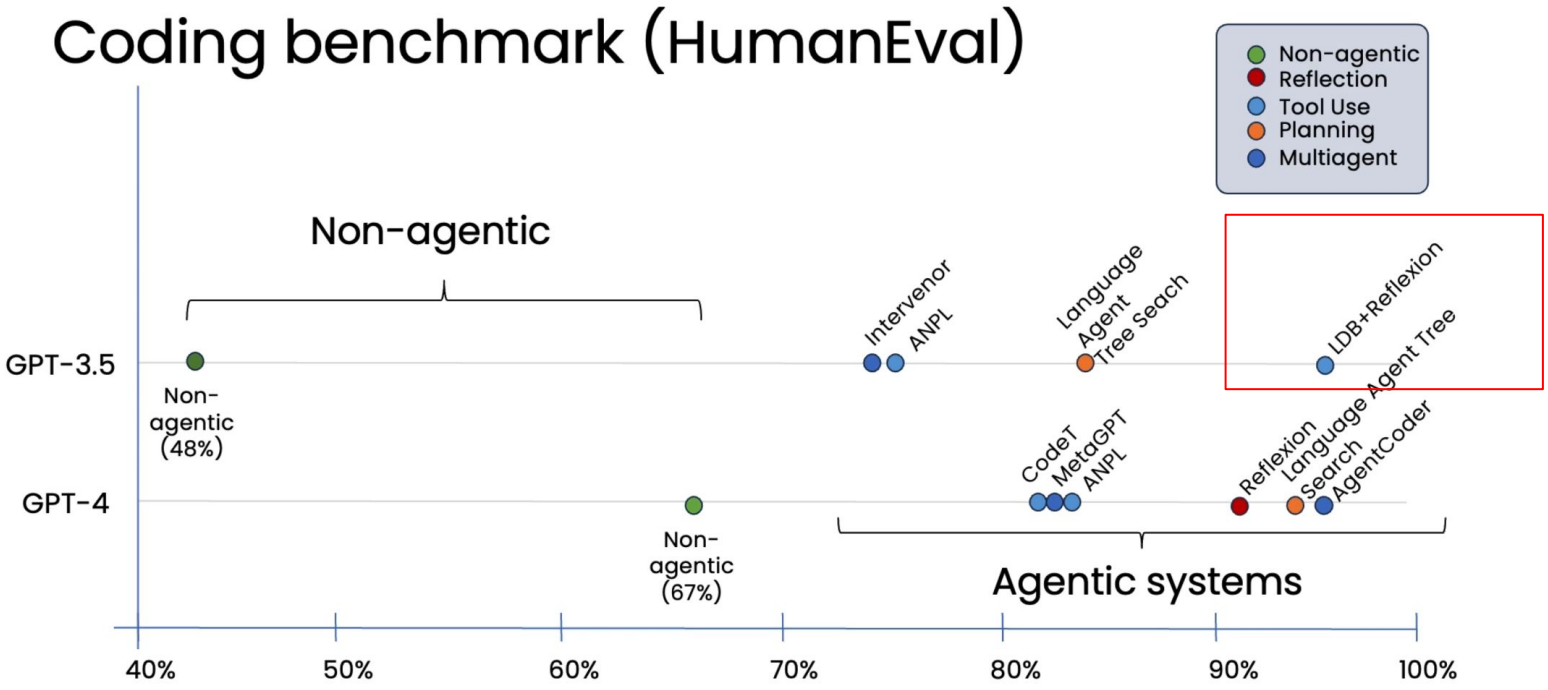

Agentic AI; Andrew Ng

"Why This Works: The Agentic Advantage"

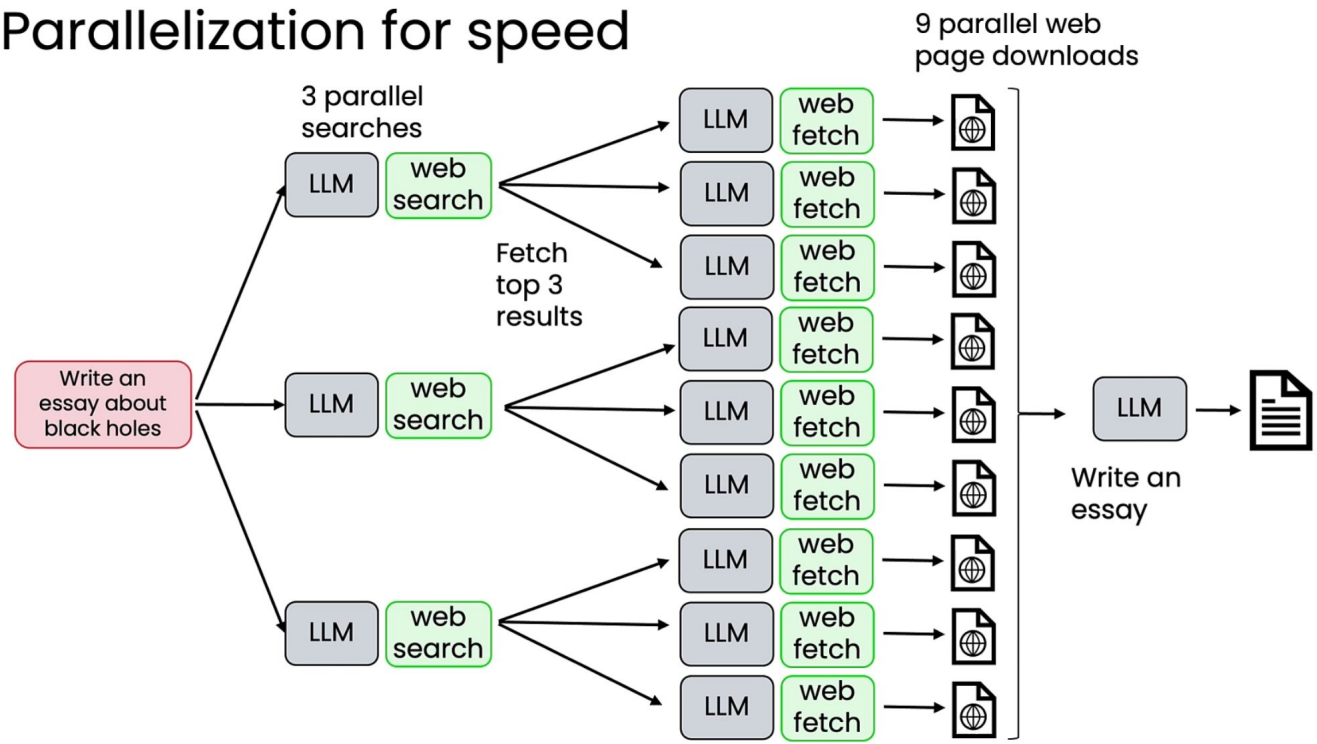

Agentic AI; Andrew Ng

"Why This Works: The Agentic Advantage"

Example: Responding to customer email

From: sjones9@email.com
Subject: Wrong item shipped

Hi, I ordered a blue KitchenPro blender (Order #8847) but received a red toaster instead.

I need the blender for my daughter's birthday party this weekend. Can you help?

Susan Jones

1. Extract key information — LLM
2. Find relevant customer records — LLM + orders database query
3. Write and send response — LLM + send email

Agentic AI; Andrew Ng

The Mental Model: Sovereignty Spectrum

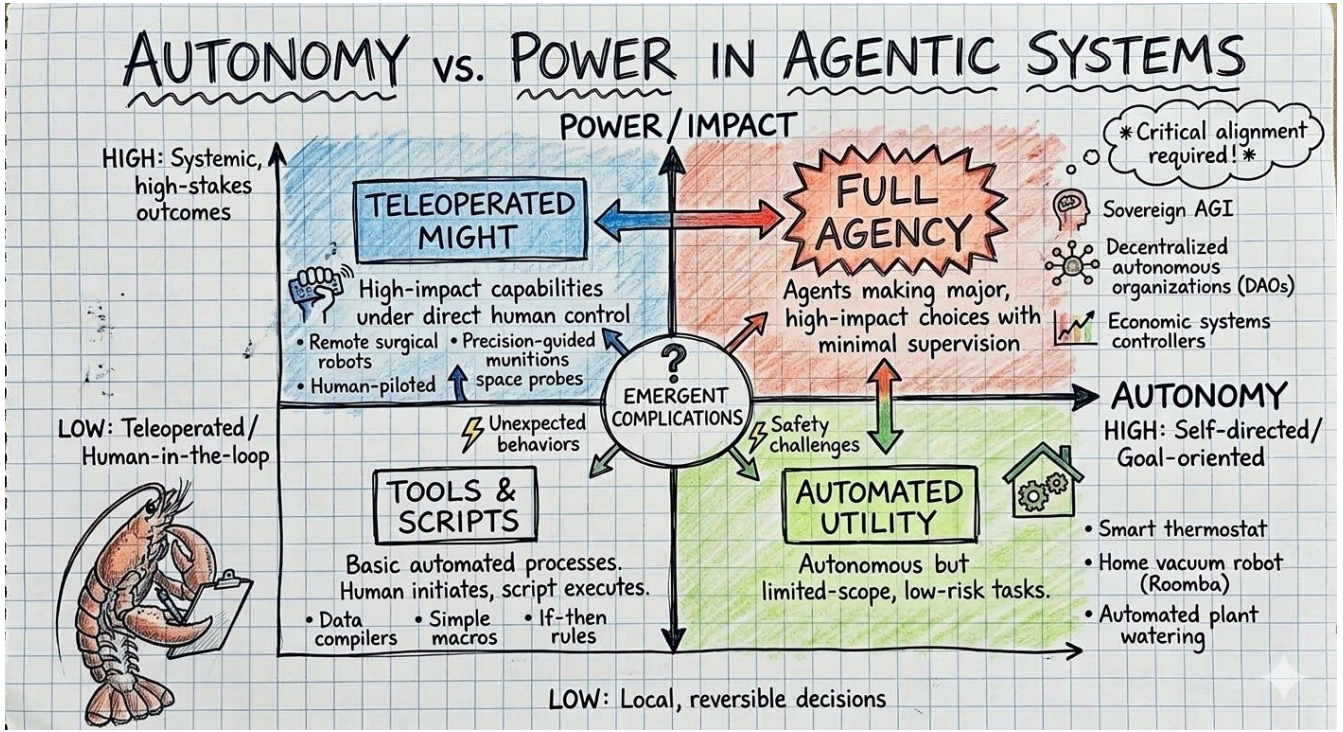

The Mental Model: Sovereignty Spectrum

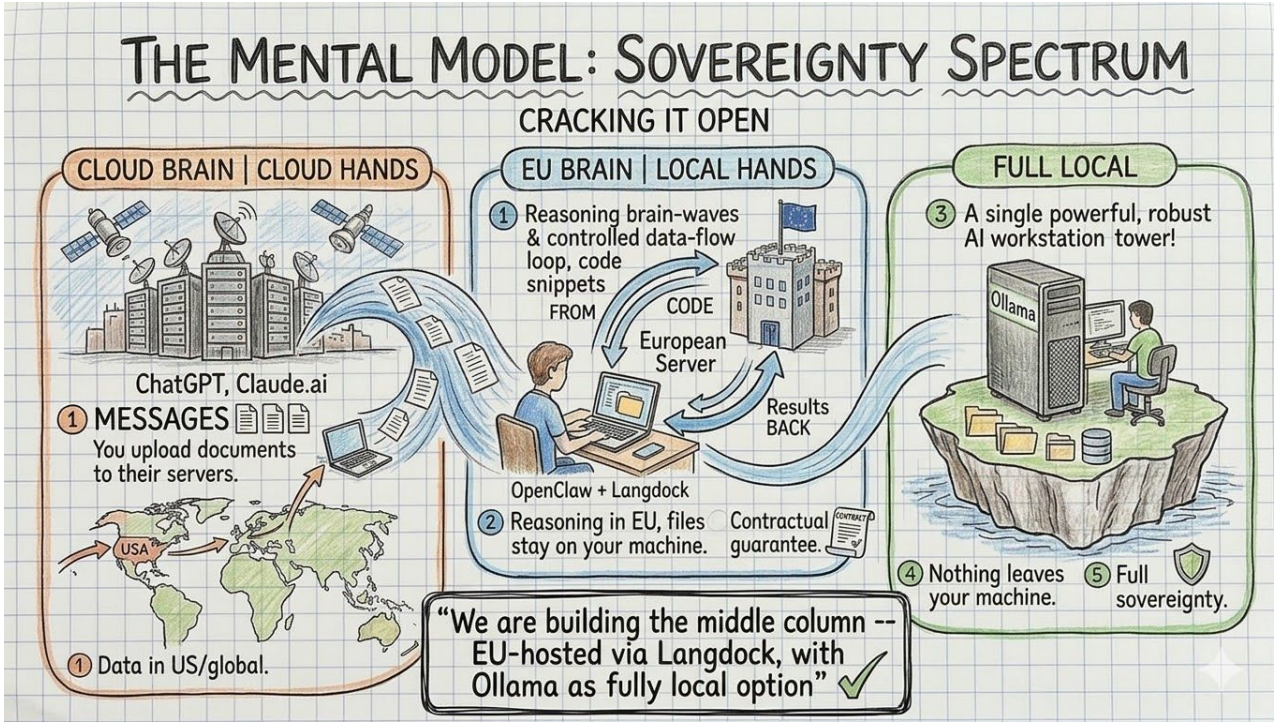

"We are building the middle column -- EU-hosted via Langdock, with Ollama as fully local option"

Why Langdock? The Sovereignty Layer

SOVEREIGNTY GUARANTEES

- ✓ German GmbH (Berlin) -- EU law jurisdiction
- ✓ Azure EU infrastructure -- data stays in EU
- ✓ ISO 27001 + SOC 2 Type II certified
- ✓ Contractual no-training guarantee
- ✓ GDPR Article 28 compliant DPA
- ✓ Zero data retention

HONEST LIMITATIONS

- When created through the platform (UI), conversations & chat histories are stored (configurable deletion)
  - API-based communication is zero-retention
- Content leaves your machine ⟶ EU servers for processing
- Multi-hop: Machine ⟶ Langdock ⟶ Azure EU ⟶ Model
- 5-10% cost premium for sovereignty layer

Langdock is not "nothing leaves your machine." It is: "files stay local, content goes to EU infrastructure with contractual guarantees, nothing trains any model." For Ollama: nothing leaves. Know your trade-off.

OpenClaw vs ChatGPT: Informed Choice, Not Sales Pitch

OPENCLAW ADVANTAGES

- AGENTS.md control -- define agent behavior in plain text
- Provider transparency -- you choose which LLM
- Audit trail -- every tool call logged to JSONL
- Workspace isolation -- files scoped per agent
- Open source (MIT) -- inspect, modify, self-host

CHATGPT ADVANTAGES

- Zero setup -- sign in and start
- Better UI -- polished, intuitive, consumer-grade
- Integrated features -- vision, code, plugins, search
- Cost clarity -- EUR 20/month, unlimited (Pro)
- Enterprise features -- SSO, admin, usage controls

"Start with ChatGPT Pro. Graduate to OpenClaw when you hit a wall: data sovereignty, proactive automation, deep customization, or internal integration."

OpenClaw Internals

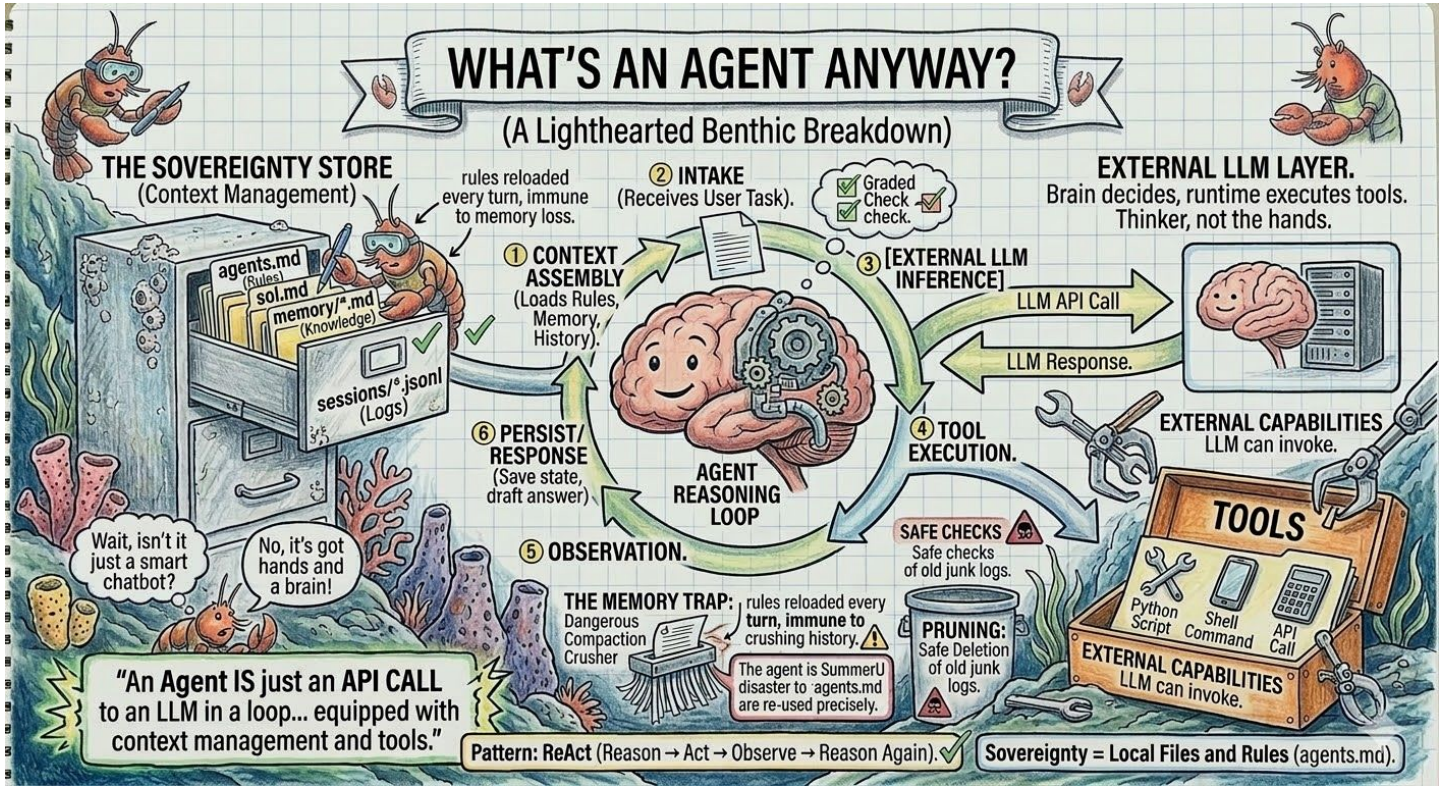

OpenClaw Internals

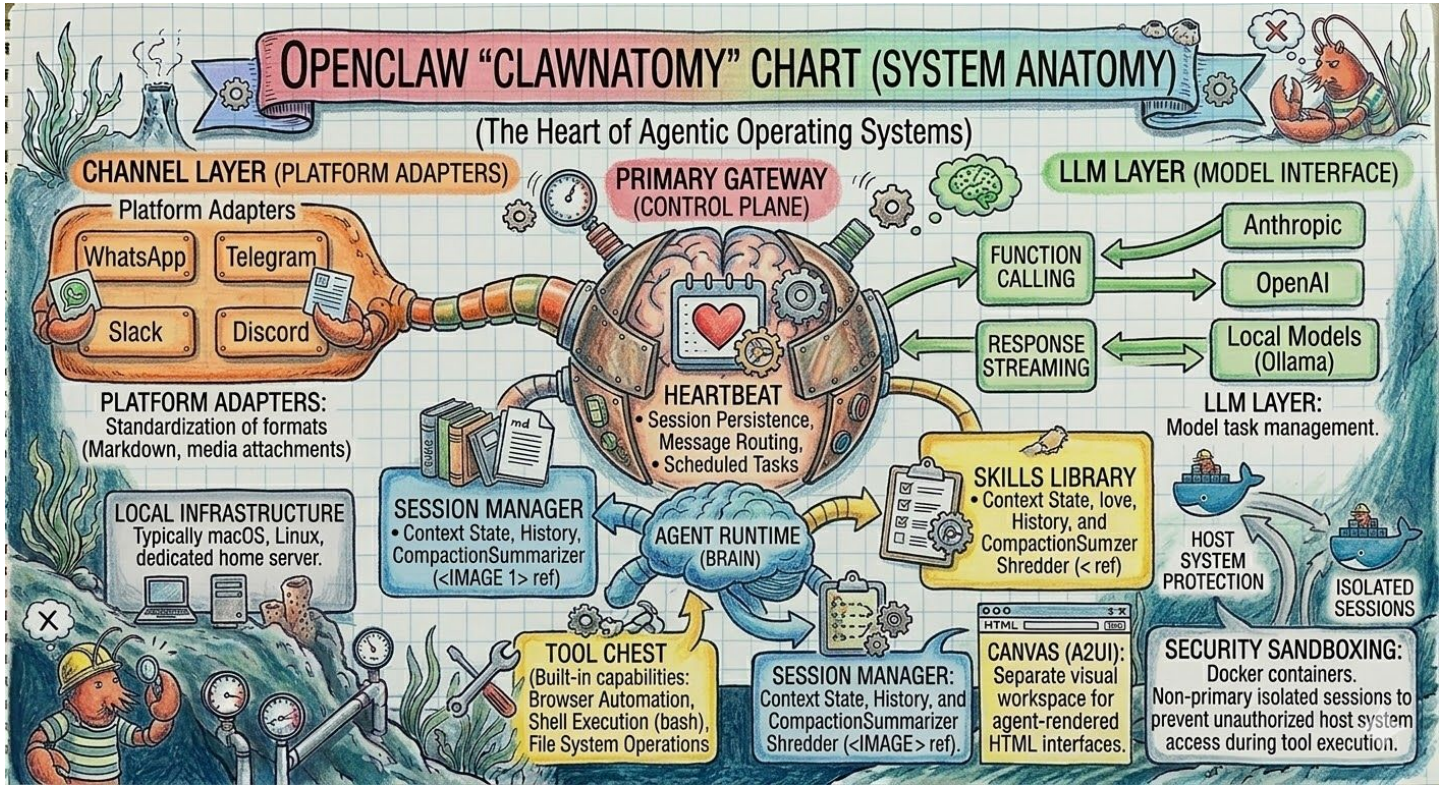

OpenClaw Internals

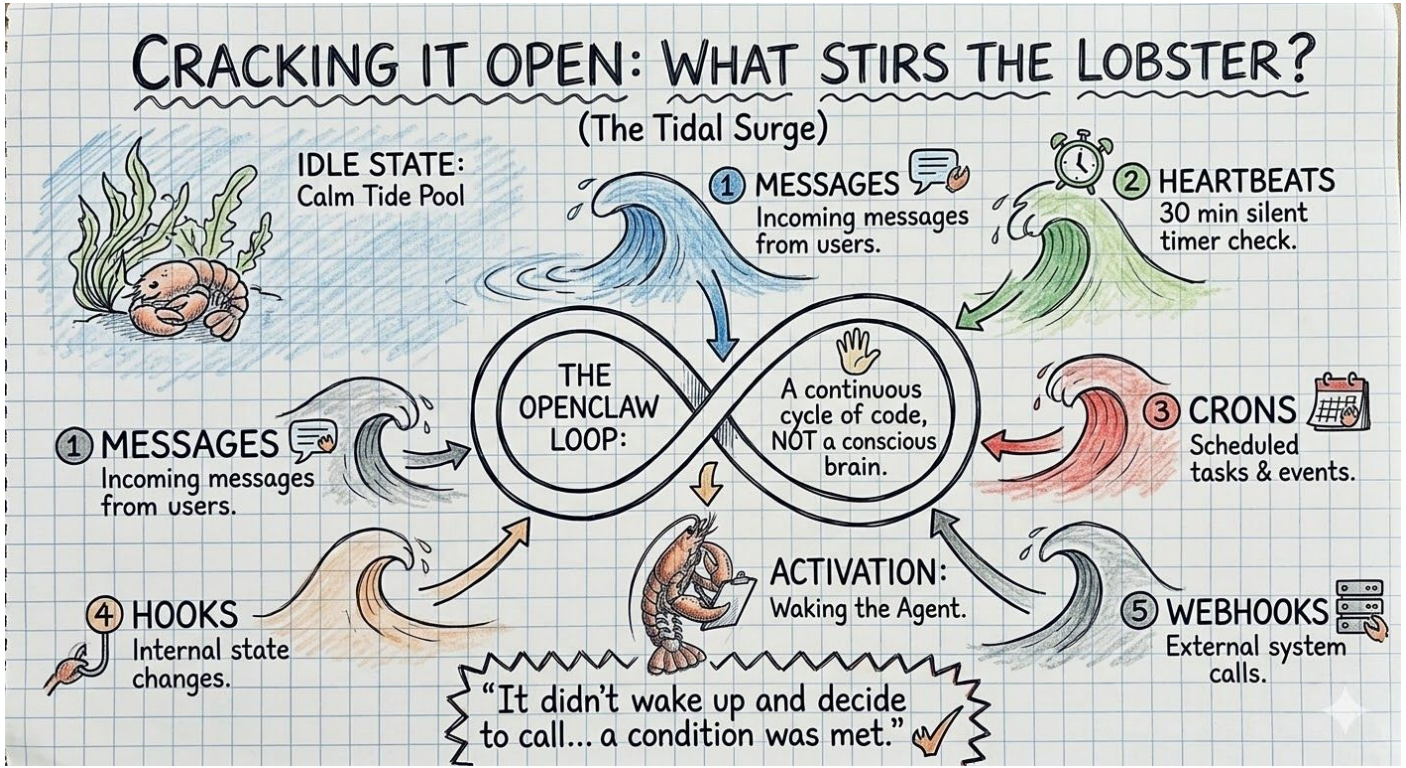

Setup Verification

Setup Checklist

- [ ] OpenClaw installed: ‘openclaw gateway’ boots up a gateway?
- [ ] Web GUI accessible: open http://127.0.0.1:18789/ in browser?
- [ ] LLM provider configured: OpenRouter API functional?
- [ ] Workspace ready: ~/openclaw-workshop/workspace/sample-docs/ has AlpenTech docs
- [ ] Test message: prompt "Hello, what can you do?" in web GUI, get a response

Troubleshooting: Check companion app --> Setup --> Troubleshooting section

CHECKPOINT

Setup Status

How's your setup?

[Thumbs Up]	[Raised Hand]	[Thumbs down]
Working	**Need help**	**Still downloading**

"We need at least 80% before we proceed"

Your First Pattern: Tool Use

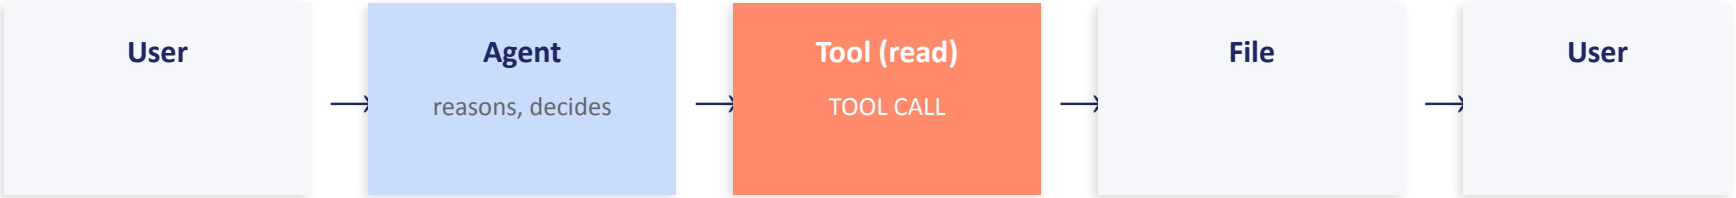

“The agent doesn't just generate text -- it takes actions.”

- Reading a file, writing output, calling an API -- that's a tool call
- The agent decides WHEN and WHICH tool to use -- that's the “agentic” part
- **In the next mission, your agent will READ a file from your workspace**
- After this mission, we'll name three more patterns. For now, focus on one: tool use.

First Tool Call

Send this to your OpenClaw:

```
1. Find the 01-consulting-agreement.txt

2. Read this document and extract: parties,
effective date, payment terms, and any risk
factors. Return as structured JSON.
```

If it works, you just built an AI agent.

"What You Just Did" -- The Four Agentic Patterns

TOOL USE

AI invokes external capabilities -- reading files, calling APIs, executing actions

PLANNING

AI breaks complex tasks into ordered steps

REFLECTION

AI evaluates and iterates on its own output

MULTI-AGENT

Multiple specialized AIs coordinate

95.1% GPT-3.5 with agentic patterns vs **67%** GPT-4 without

"What Could Go Wrong?" -- Security Checkpoint

RISK	MITIGATION
RISK 1 Uncontrolled file access	MITIGATION Allow only the document directory
RISK 2 Prompt injection via document	MITIGATION Structured output schema constraints response
RISK 3 Messaging channel compromise	MITIGATION Workshop-scoped, not internet-exposed

"Sovereign deployment = architectural security, not prompt-level security"

Break -- 10 Minutes

Stretch, refill, check your extraction results.
When we return: building a full document intelligence pipeline.

Document Intelligence Pipeline

"From one document to a pipeline"

Contract
Invoice | Financial Report | HR Policy | Research Paper

The Extraction Schema

```
{ "extraction": {
    "document_id": "string",
    "document_type": "contract|invoice|report|...",
    "entities": {
      "organizations": [...],
      "people": [...],
      "dates": [...],
      "monetary_values": [...]
    },
    "key_facts": [...],
    "structured_data": { /* type-specific */ },
    "summary": "string"
  }
}
```

Schema = Quality Guarantee

- Standard schema for all document types
- Entities: organizations, people, dates, monetary values
- Key facts: extracted as structured bullet points
- Type-specific structured_data adapts to each document
- Makes output comparable and machine-parseable

Mission 2: What Your Agent Will Do

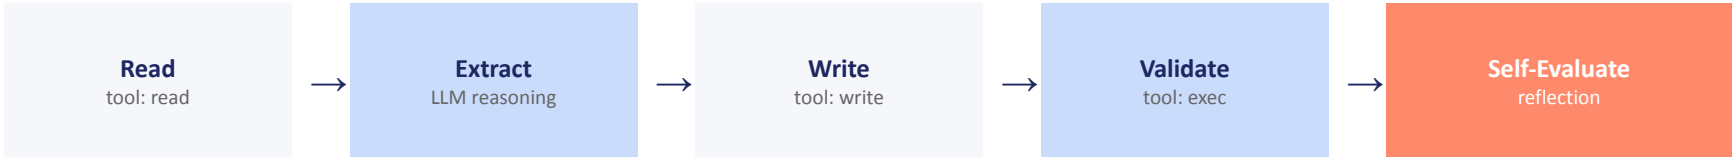

"The sequence may vary. The agent decides the order -- that's planning in action."

Watch for three tool types:

- **read** -- file access (reading the document from workspace)
- **write** -- file output (writing extraction results to JSON)
- **exec** -- validation (running JSON validation on its own output)

Open companion app ⟶ Missions ⟶ Mission 2 for full briefing and gold standards

Process the Full Corpus

- Open companion app --> Missions --> Mission 2

- Process all five sample documents through your OpenClaw pipeline
- Use the extraction schema to constrain output
- Compare results across document types
- Time: ~25 minutes hands-on
- Gold standard results available for comparison

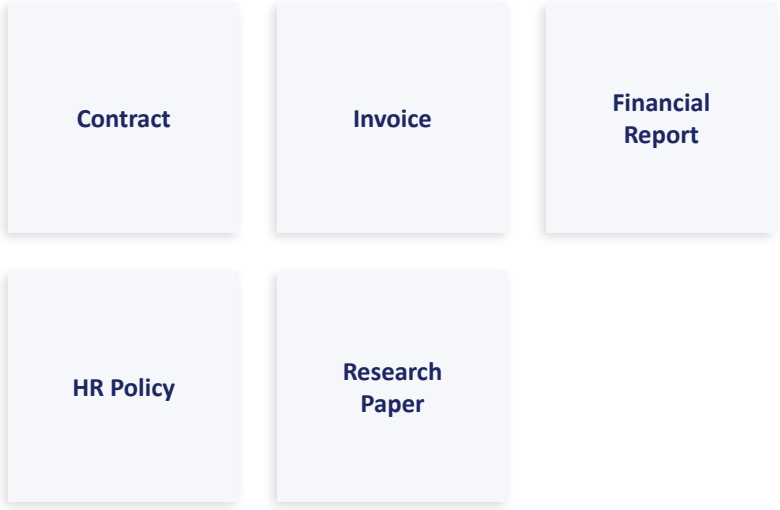

DISCUSSION

"What Just Happened?"

- You built a sovereign document intelligence pipeline
- Three agentic patterns in action: Tool Use, Planning, Reflection
- Can you verify WHERE your data was processed?

- The sovereignty spectrum: what can you tell your compliance team?

Discussion:

"Where would this be useful in your work? Which bottlenecks can you identify which may benefit from an agentic system"

ClawBio: The Wave Goes Deeper

github.com/ClawBio/ClawBio

- 7 production skills, 57 tests, MIT licensed
- Domain-specific OpenClaw fork for computational biology
- Pharmacogenomics analysis in under 1 second -- locally
- Community PR within 24 hours of demo

"80% of apps are going away"

-- Peter Steinberger

- Pattern spreading: domain-specific forks emerging everywhere
- International adoption including China
- The "agents replace apps" thesis playing out in specialized fields

Preview: Day 2

File-as-Interface	Control AI behavior with text files
Skills & MCP	Plug in new capabilities
Multi-Agent	When one isn't enough
EU AI Act Compliance	From demo to deployment

Day 1 Complete

- ✓ Working OpenClaw + Ollama setup
- ✓ Processed documents with structured extraction
- ✓ Mental model of agentic AI patterns
- ✓ Understanding of the sovereignty spectrum

Tomorrow: 16:00 CET, "Making It Yours"

Questions? Chat, email, or stay on the call

The Sovereignty Question

Scenario 1: Cloud	Scenario 2: EU-Hosted	Scenario 3: Full Local
"We uploaded the contract to US servers"	*"Files stayed local, content processed on EU infrastructure with guarantees"*	*"Nothing left the machine, here's the log"*
Compliance nightmare	**Defensible**	**Bulletproof**

EU AI Act (effective August 2026): transparency, auditability, data governance
"Sovereignty isn't just about privacy -- it's becoming a compliance requirement"

"Here's What We'll Do Instead"

THE PROBLEM

- OpenClaw: uncontrolled exposure
- Internet-facing, no sandboxing
- Marketplace with malicious skills
- No audit trail

OUR APPROACH

- Same capability: AI reads, processes, acts
- Containerized, local-only, no marketplace
- Files stay on your machine. Content processed on EU infrastructure (Langdock) or fully local (Ollama).
- Full audit trail

"Same power, different architecture"

Sovereign Agentic Systems: Europe-Centric Agentic Workflows for Sensitive Documents; Day 2

Shaping the Future of AI

Dejan Đukić
Senior AI/ML Engineer, TetraScience

18.03.2026

Day 1 Recap (60-Second Version)

Setup Working

OpenClaw + Langdock (EU) or Ollama (local)

Documents Processed

Five types extracted to JSON

Patterns Understood

Tool Use, Planning, Reflection

Sovereignty Spectrum

Cloud vs EU vs Local

Today: How to customize, extend, build skills

Day 1 Recap

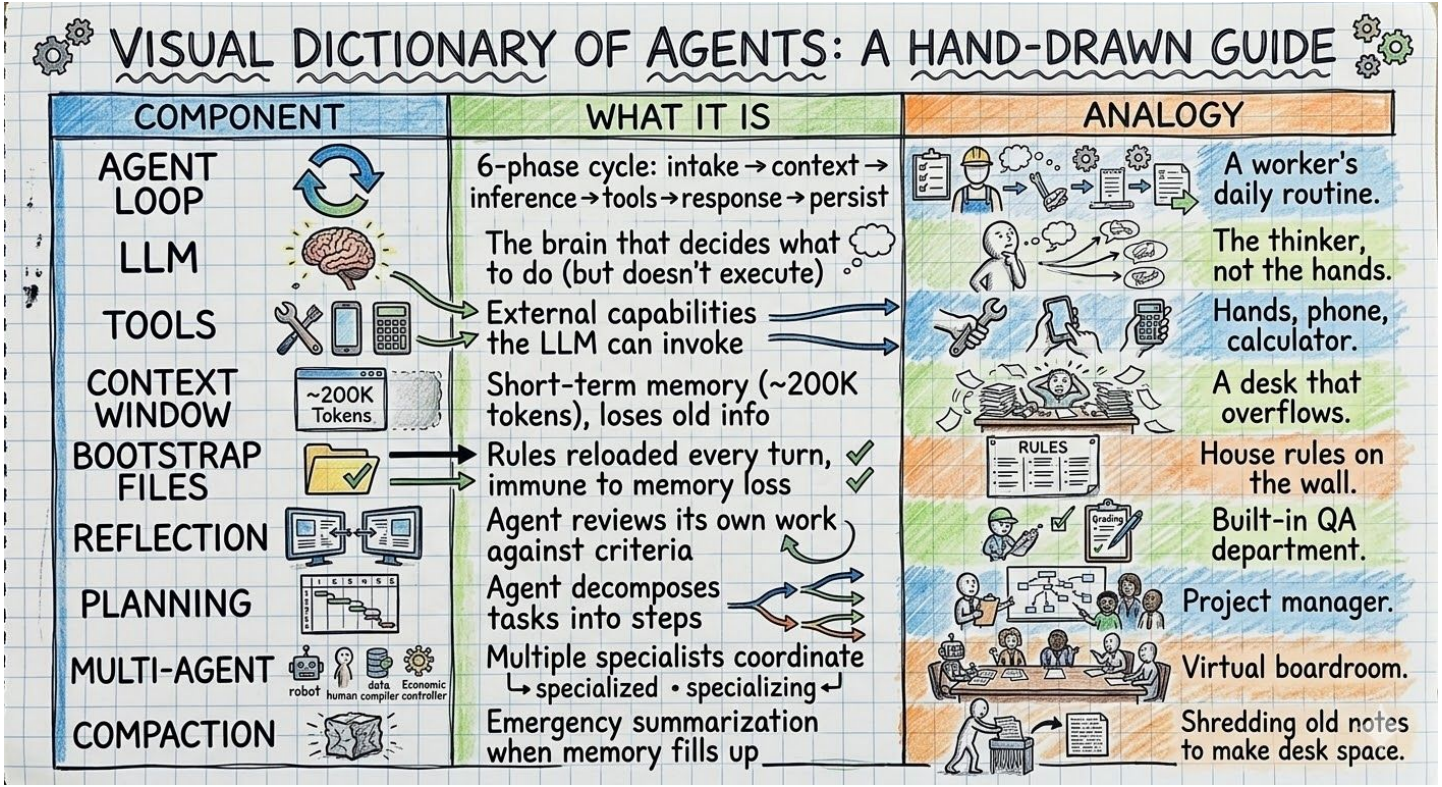
VISUAL DICTIONARY OF AGENTS: A HAND-DRAWN GUIDE
COMPONENT
WHAT IT IS
ANALOGY
AGENT LOOP
6-phase cycle: intake → context → inference → tools → response → persist
A worker's daily routine.
LLM
The brain that decides what to do (but doesn't execute)
The thinker, not the hands.
TOOLS
External capabilities the LLM can invoke
Hands, phone, calculator.
CONTEXT WINDOW
~200K Tokens
Short-term memory (~200K tokens), loses old info
A desk that overflows.
BOOTSTRAP FILES
Rules reloaded every turn, immune to memory loss
RULES
House rules on the wall.
REFLECTION
Agent reviews its own work against criteria
Grading
Built-in QA department.
PLANNING
Agent decomposes tasks into steps
Project manager.
MULTI-AGENT
robot
human
data compiler
Economic controller
Multiple specialists coordinate
↳ specialized • specializing ↵
Virtual boardroom.
COMPACTION
Emergency summarization when memory fills up
Shredding old notes to make desk space.

Skills

- Agent Skills are a lightweight, open format for extending AI agent capabilities.
- A **skill** is a **folder of organized files** consisting of instructions, scripts, assets and resources that agents can discover to perform a specific task accurately.

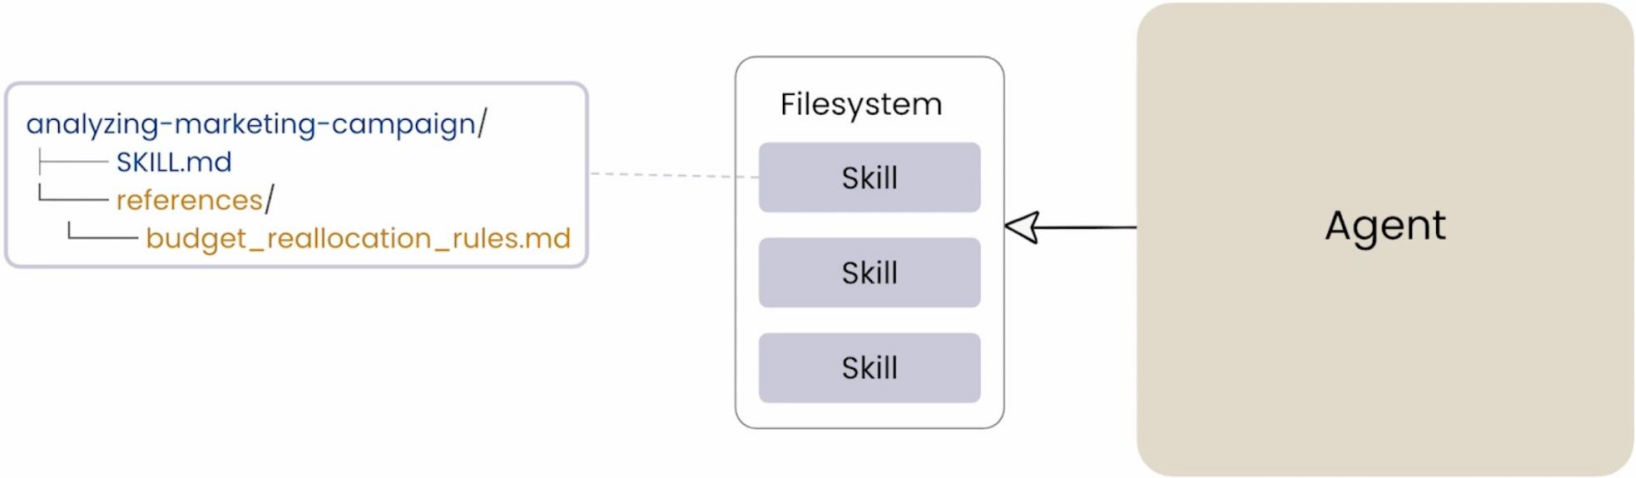

Skills

How we used to think about agents

General-Purpose Agent:
Code is the universal interface

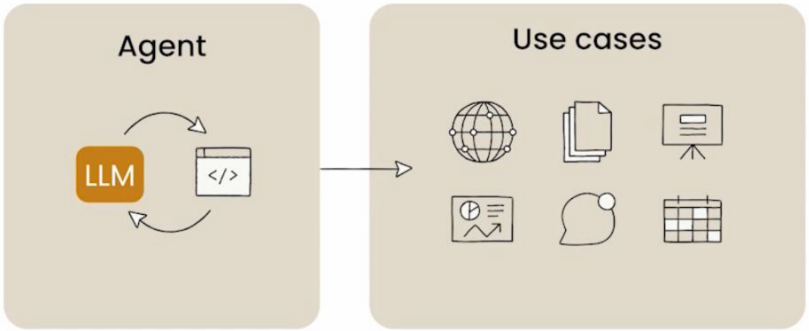

- Simple Scaffolding: **bash** and **filesystem**.
- But they need **context** and **domain expertise** to do the job **reliably.**

Skills

Use Cases

Domain Expertise

- Brand guidelines and templates
- Legal review processes
- Data analysis methodologies

Repeatable Workflow

- Weekly marketing campaign review
- Customer call prep workflow
- Quarterly business review

New Capabilities

- Creating presentations
- Generating Excel sheets or PDF reports
- Building MCP servers

Without Skills

- Describe your instructions and requirements every time
- Bundle all your references and supporting files every time
- Ensure the workflow or outputs are always consistent

Skills

analyzing-marketing-campaign/SKILL.md

YAML Frontmatter

```
---
name: analyzing-marketing-campaign
description: Analyze weekly marketing campaign performance data
across channels. Use when analyzing multi-channel digital
marketing data to calculate funnel metrics and compare to
benchmarks, compute cost and revenue efficiency metrics, or get
budget reallocation recommendations based on performance rules.
---
```

Metadata: always loaded

Markdown

```
## Input

Excel/CSV with columns: **Date, Campaign_Name, Channel, Segment, Impressions,
Clicks, Conversions, Spend, Revenue, Orders**

## Funnel Metrics (per channel)

- **CTR** = Clicks / Impressions × 100
- **CVR** = Conversions / Clicks × 100

## Efficiency Metrics (per channel)

- **ROAS** = Revenue / Spend (target: ≥4.0x)
- **CPA** = Spend / Conversions (max: $50)
- **Net Profit** = Revenue - (Spend + Orders×$8 + Revenue×35%)

## Output Tables

**Funnel:** | Channel | CTR Actual | CTR Benchmark | CTR Diff | CVR Actual | CVR
Benchmark | CVR Diff |
**Efficiency:**  Channel | ROAS | Status | CPA | Status | Net Profit | Status |
Include brief interpretation and recommendations per channel

## Budget Reallocation

If requested, see `references/budget_reallocation_rules.md` for decision
framework.
```

Instructions: loaded when triggered

Resources: loaded as needed

```
## Weekly Optimization Framework

### Rule 1: Performance-Based Decision Framework
Use the following criteria to determine budget adjustments for each marketing
channel on a weekly basis.

**Increase Budget by 10-15%**

A channel qualifies for budget increase when it meets all of the following
conditions:
- ROAS exceeds the target threshold by 15% or more

... and so on ...
```

Sovereign control

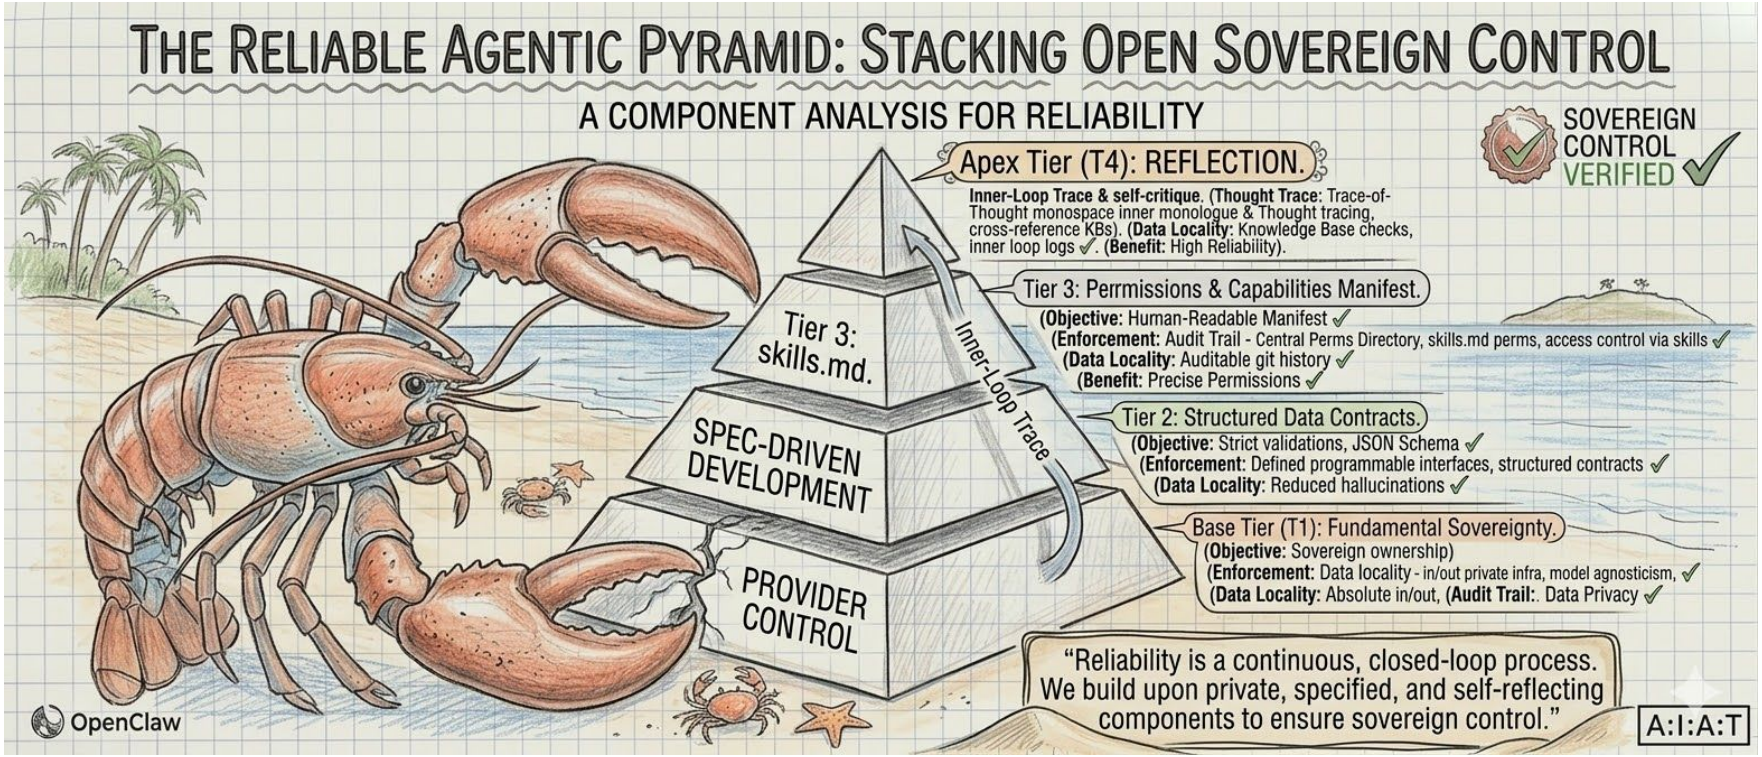

Sovereign control elements

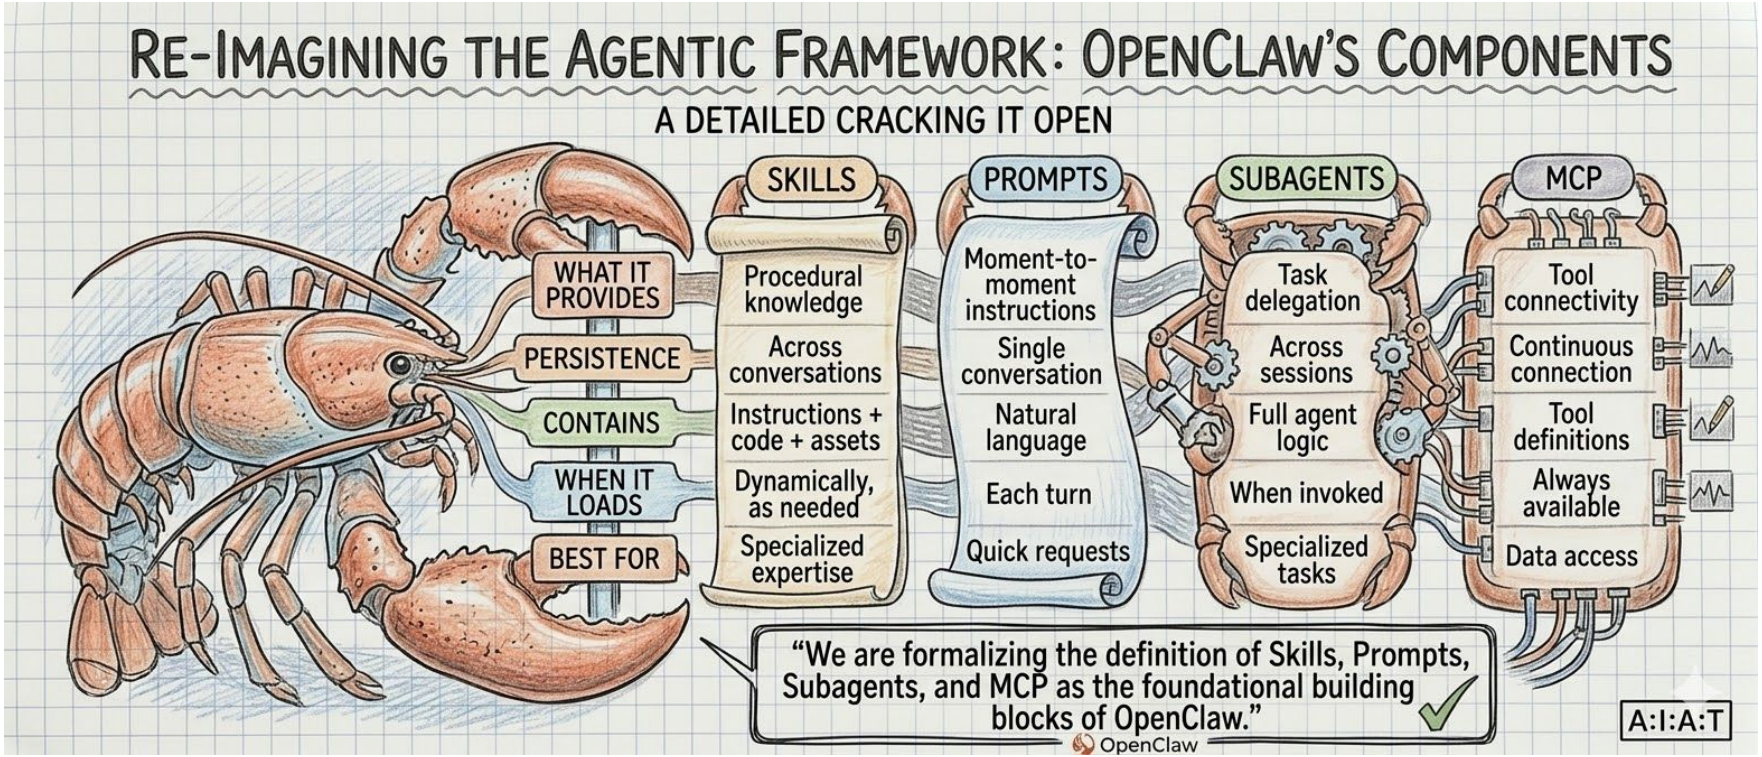

Agenda: Day 2

16:00-16:10	**Recap + Mission 2 Briefing**	finish what we started yesterday
16:10-16:25	**MISSION 2: The Audit**	process 5 docs, validate JSON, reflect (hands-on)
16:25-16:40	**File-as-Interface + Skills**	AGENTS.md, SKILL.md, OpenClaw customization
16:40-16:55	**MISSION 3: The Control File**	edit AGENTS.md, build a custom skill (hands-on)
16:55-17:05	**Break**	stretch, refill
17:05-17:25	**Multi-Agent Patterns + ClawNet**	when one agent isn't enough
17:25-18:00	**MISSION 4 + EU AI Act + Closing**	agent teams (hands-on), audit dashboard, business case

MISSION

The Audit

Process All 5 Documents

Agreement, invoice, financial report, HR policy, research paper

Write Structured JSON

Agent writes extraction-results.json to outputs/

Validate Programmatically

node -e JSON.parse(...) -- not just 'looks right'

Agent Self-Evaluation

Confidence scores + reflection report written to outputs/

PII Detection

Find hidden SV-Nr, passport numbers, privileged content

Carried from Day 1 | Time: ~15 min | Patterns: Tool Use + Planning + Reflection

AGENTS.md

```
# AGENTS.md

You are a document intelligence agent.
You process legal and financial documents.

## Behavior Rules
- Always extract entities, dates, and monetary values
- Never store or transmit document contents externally
- Respond in the user's language
- Flag any clauses related to liability or termination

## Output Format
- Always return structured JSON
- Include confidence scores for each extraction
```

No code
No YAML
No JSON config syntax

Just natural language.

See also, CLAUDE.md, SOUL.md

The File-as-Interface Paradigm

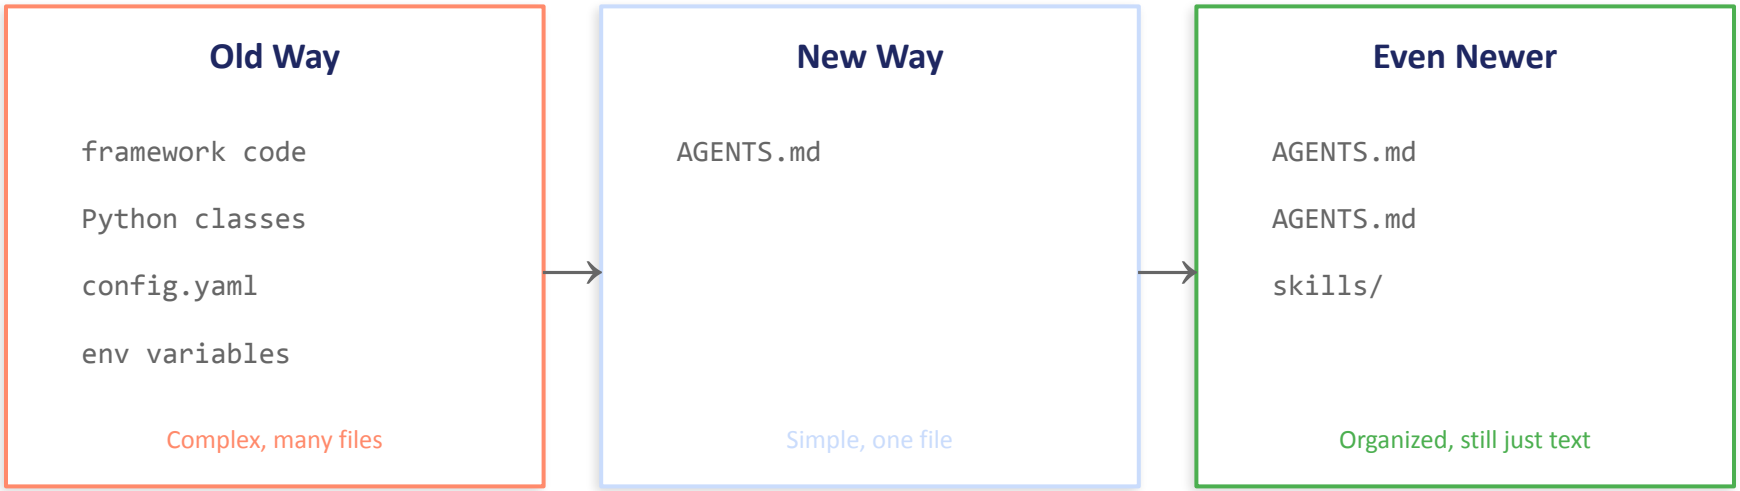

AGENTS.md: Global behavior -- personality, constraints, priorities
AGENTS.md: Role-specific instructions -- "you are a legal document analyst"
System prompt: Per-session overrides

This is why you don't need to be a programmer to orchestrate AI

Changing Behavior with a File Edit

LIVE DEMO

Before:

[SCREENSHOT: Baseline agent behavior -- English, general extraction]

After:

[SCREENSHOT: German response, risk factors only]

```
+ "When processing documents, focus exclusively on financial risk factors

  and respond in German."
```

"Configuration as conversation" -- no restart, no recompile

AGENTS.md: Role-Based Configuration

Legal Analyst

```
You review contracts for
risk.
Focus on liability,
termination,
and regulatory compliance.
```

AGENTS.md

Financial Auditor

```
You extract KPIs and flag
anomalies. Compare against
industry benchmarks.
```

AGENTS.md

Research Summarizer

```
You read papers and
extract
methodology, findings, and
limitations.
```

AGENTS.md

"Think of AGENTS.md as a job description for your AI"

Planning & Task Decomposition

[OK]

Complex Task Decompose

Each step: can an LLM or tool do this?

[!!]

Complex Task Into Steps

If no,
decompose further

[Teams]

Agent Plans Dynamically

Planning pattern:
agent decides steps

If no, "How would YOU do this manually? Those steps become the agent's steps."
"A good employee doesn't need step-by-step instructions for every task"

Skills & MCP

"How Your Agent Learns New Tricks"

Text files control behavior. Skills add capabilities.

How Skills Work

Models = Processors | Agent Runtimes = Operating Systems | Skills = Applications

"You're not programming the agent. You're giving it tools and letting it decide."

MCP: "USB-C for AI Tools"

USB-C

One port, many devices

Keyboard | Monitor | Storage | Phone

MCP (Model Context Protocol)

One protocol, many tools

Files | **Database** | **Web API** | **Docs**

- MCP = Model Context Protocol (Anthropic, open standard)
- Standardized way for AI to connect to external tools
- Growing ecosystem: file systems, databases, APIs, document tools
- PageIndex, Slack, GitHub, custom tools -- all through MCP
- "You don't need to learn a new API for each tool"

Skill Anatomy: What's Inside

```
# Skill: Compliance Checker

## Metadata
name: compliance-checker
description: Checks documents for regulatory compliance
version: 1.0

## When to Use
Activate when the user sends a legal or regulatory document
and asks about compliance, risk, or violations.

## What to Do
1. Extract all regulatory references
2. Check against known frameworks (GDPR, EU AI Act)
3. Flag potential non-compliance issues
```

1. METADATA

Name, description, version

2. WHEN TO USE

Trigger conditions in natural language

3. WHAT TO DO

Capability definition -- instructions, not code

MISSION

The Control File

Edit AGENTS.md

Add one line, observe immediate behavior change

No Restart Needed

OpenClaw reads AGENTS.md fresh every session

Install Skill-Builder

npx clawhub@latest install skill-builder

Create Custom Skill

Package your workflow as a reusable /command

Test Your Skill

Invoke via slash command -- consistent, repeatable output

Time: ~20 min | Patterns: File-as-Interface + Skills | Change the file, change the agent

MISSION

The Control File + Build a Skill

Compliance Checker

Flags clauses that might violate regulations

Executive Summary

Produces a one-page brief from any document

Comparison Extractor

Compares two documents, highlights differences

Regulatory Watch

Monitors folder for regulatory changes

Meeting Actions

Extracts action items, owners, deadlines

Edit AGENTS.md then create a skill | Time: ~20 min | Use OpenClaw to help write the skill

CHECKPOINT

Skill Building Results

What did you build?

- Did the agent use it correctly? When you expected?
- What surprised you?
- "The agent decided when to use your skill -- you didn't code that logic"
- Common gotcha: vague trigger conditions lead to inconsistent activation

Break -- 10 Minutes

When we return: what happens when one agent isn't enough.
The fourth and final agentic pattern.

"When One Agent Isn't Enough"

[OK]

Single Agent
+ Simple Task

Manageable

[!!]

Single Agent
+ Complex Task

Context rot:
agent 'forgets'

[Teams]

Specialized Agents
Divide Work

Efficient
delegation

Context rot: the agent "forgets" earlier information as context fills up
"This is why large organizations need multi-agent systems"

Multi-Agent Patterns

SUPERVISOR

Boss agent delegates tasks, collects results

Most common

ROUTER

Incoming work sent to the right specialist

Classification-first

PIPELINE

Sequential processing -- each agent does one step

Document assembly lines

SWARM

Autonomous peer-to-peer communication

Advanced, emerging

"Most of you will need Supervisor or Router. Swarms are for the brave."

Ambient Intelligence Done Right

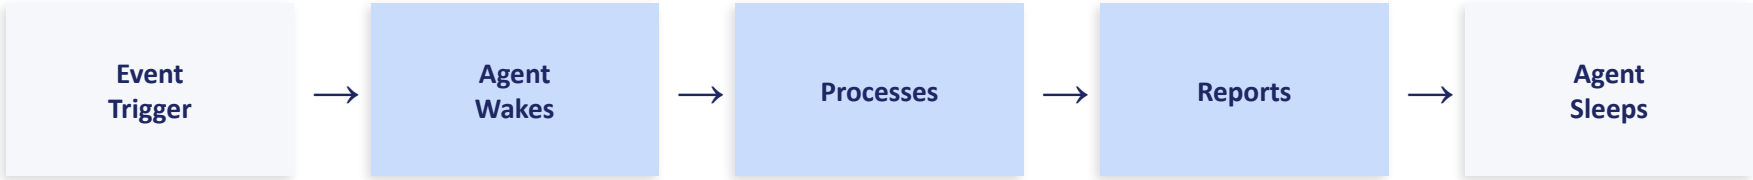

Five Input Types:

The 3 AM phone call:
Your agent calls at 3 AM -- it's not sentient, it's a cron job that detected something matching your alert criteria.
"Your agent should be like a good assistant -- available when needed, not hovering."

Multi-Agent in Action

LIVE DEMO

- Supervisor agent receives a mixed document folder
- Classifies each document, routes to specialist agent
- Legal agent handles contracts, financial agent handles reports
- Supervisor collects results, produces synthesis
- "Each agent has its own AGENTS.md -- specialized behavior from text files"

MISSION

Agent Teams

Orchestrator (EU)

Langdock -- routes tasks based on data sensitivity

Analyst (EU)

Langdock -- handles PII, contracts, sensitive docs

Researcher (US)

OpenRouter -- web research, public information only

Test Multi-Step Routing

EU AI Act research then contract compliance analysis

Defense in Depth

5 safety layers -- what if routing fails?

Time: ~20 min | Patterns: Multi-Agent + Tool Use | Never rely on a single control

The Four Patterns Complete

TOOL USE

Day 1

AI invokes external capabilities

PLANNING

Day 1

AI breaks tasks into ordered steps

REFLECTION

Day 1

AI evaluates and iterates on output

MULTI-AGENT

Day 2

Multiple specialized AIs coordinate

"These patterns are the transferable knowledge. They apply to any framework, any tool."

Context Rot: Pruning vs Compaction

PRUNING

Safe: removes used tool outputs
Like clearing junk mail
Core conversation preserved

[SAFE]

COMPACTION

Dangerous: summarizes entire history
Safety rules can vanish
Summer Yu's disaster was this

[DANGEROUS]

THE FIX:

- **Critical rules go in bootstrap files (AGENTS.md), not in chat**
- Split long tasks into fresh sessions
- Use structured handoff documents between sessions
- Dedicated QA agent that checks for contradictions

Agent Memory: From Files to Knowledge Graphs

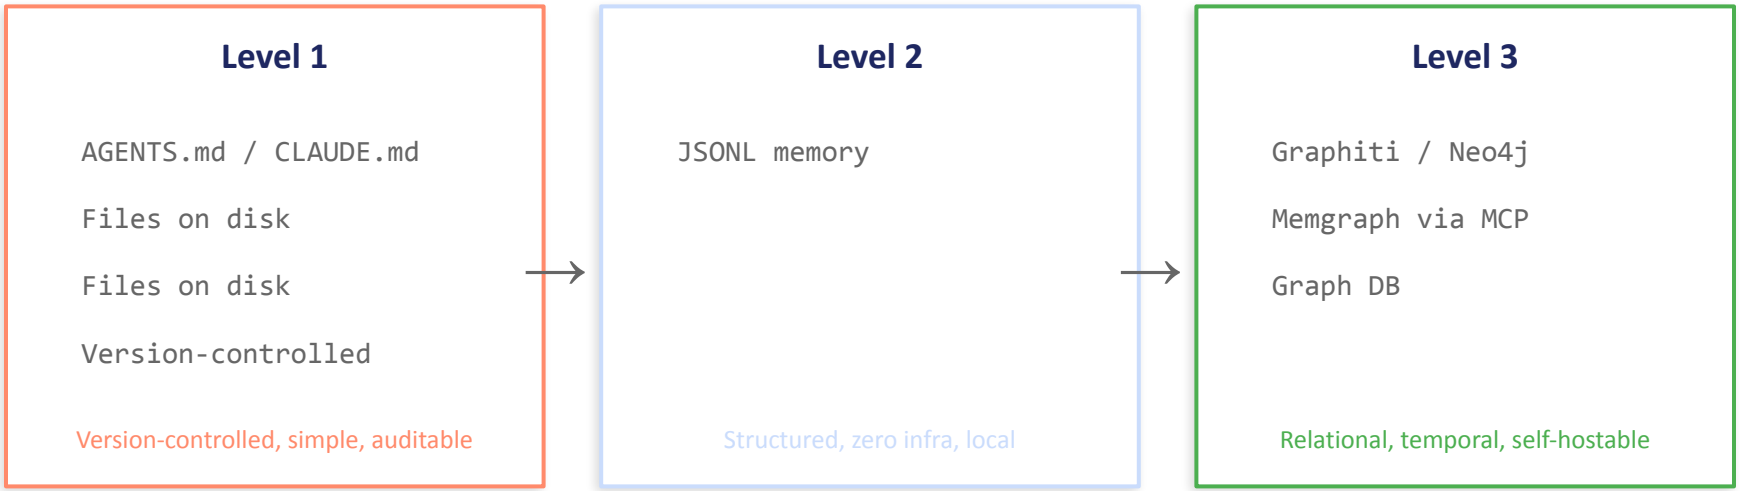

Level 1: AGENTS.md IS your memory -- version-controlled, auditable, sovereign
Level 2: MCP tools store entities in local JSONL -- add in 2 min, zero infra
Level 3: Self-hosted graph DB -- tracks what changed and when
"RAG finds what sounds right. Memory tracks what changed over time."

Human-in-the-Loop: Gates and Review Points

[Green] = automated [Blue] = human gate

- Fully autonomous agents are rarely appropriate in production
- Gates: points where a human reviews and approves before proceeding
- Audit trails: every decision logged, every tool call recorded
- "The goal isn't replacing humans -- it's augmenting them with transparency"

EU AI Act: "GDPR for AI"

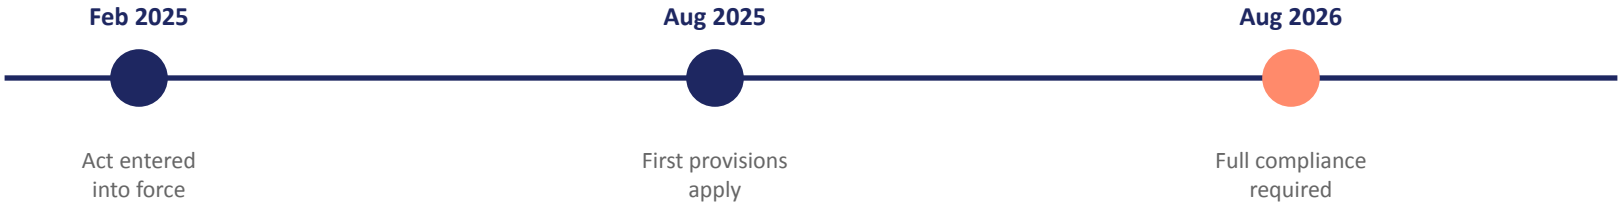

Transparency
Users must know they're interacting with AI

Auditability
Decisions must be traceable and explainable

Data Governance
Processing data must be documented

"Think of this like GDPR but for artificial intelligence." -- IBM Technology

Langdock EU hosting + sovereign audit dashboard + AGENTS.md = compliance-ready

The Decision Framework

What are you trying to do with AI?

Process Sensitive Documents	Build Internal Tools	Customer-Facing AI
Full local (Ollama) or EU-hosted (Langdock)	Start with single agent + skills, grow to multi-agent	Maximum hardening, human gates, audit all

When NOT to use agentic AI: Simple queries, non-sensitive data, existing tools work fine

"The best AI architecture is the simplest one that solves your problem."

The Sovereignty Spectrum (Final Reference)

	Cloud (ChatGPT/Claude)	EU-Hosted (Langdock)	Full Local (Ollama)
Data location	US/Global	EU (contractual)	Your machine
Capability	Frontier models	Frontier models	Smaller (80-90%)
Cost	Per-token	~2-3 EUR/workshop	Free (HW cost)
Compliance	Difficult	Defensible	Bulletproof
Setup	Minutes	Hours	Hours + hardware
Best for	Non-sensitive exploration	Production with compliance	Maximum sovereignty

AI Slop vs Cognitive Sovereignty

AI SLOP

- "Looks good to me" without reading
- Shipping first-draft AI output as final
- Feeling confident while understanding less
- AI vocabulary leaking into your writing

COGNITIVE SOVEREIGNTY

- Use AI as a sparring partner, not ghostwriter
- Have AI ask YOU questions (Jeremy Utley)
- Check biases -- ask AI to argue the opposite
- Never skip the reputation test

"Same tool. Different mindset."

Resources & Further Learning

LEARN MORE

- Companion website: https://agentic-workshop-aiat.vercel.app/
- Andrew Ng's Agentic AI Course (DeepLearning.AI)
- Anthropic: "How we built our multi-agent research system"
- Anthropic: "Don't Build Agents, Build Skills Instead" (866K views)
- LangChain: "Choosing the Right Multi-Agent Architecture"

TOOLS

- OpenClaw documentation
- MCP specification & ecosystem
- PageIndex (VectifyAI) -- vectorless document reasoning
- IronClaw -- security-first Rust agent
  GLiNER2 -- 205M param entity recognition, CPU
  Benchmark results in companion app Experiments section

VIDEO REFERENCES

- Niklas Steenfatt: German Ollama guide (61.7K views)
- Fahd Mirza: Local Ollama deployment (175K views)
- Brian Casel: Multi-agent team architecture (510K views)
- VelvetShark: 50-day honest assessment (183K views)

COMMUNITY

- AI Factory Austria -- upcoming workshops and events
- AGENTS.md specification -- the emerging standard
- ClawBio (github.com/ClawBio/ClawBio)

Questions & Discussion

"What will you build first?"

"What would you need to convince your team to try this?"

"What's still unclear?"

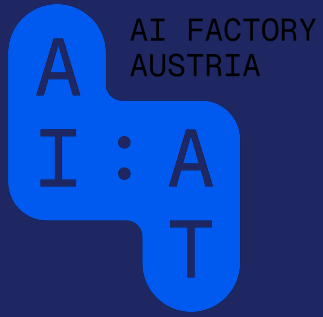

Cogito, ergo sum.

“I think, therefore I am.” -- Rene Descartes, 1637

How much are we willing to pay
to liberate our time —
for our families, our hobbies, our lives?

And if AI does your thinking...
what’s left of you?

AI Factory Austria -- AI:AT

Training & Skills Development

• AI Workshops • Courses • Webinars

ai-at.eu/trainings

Scan to explore
our events

Funded by

under discussion with

EuroHPC
Joint Undertaking

Funded by
the European Union

Federal Ministry
Innovation, Mobility
and Infrastructure
Republic of Austria

FFG
Promoting Innovation.

AI Factory Austria AI:AT has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101253078. The JU receives support from the Horizon Europe Programm of the European Union and Austria (BMIMI / FFG).