

INTRAFIND

A prompt in a multi-turn chat is not an agent!

SCCON 2026 · Daniel Manzke, Head of Engineering



THE PROBLEM

Agent washing: the label outran the technology



~130

of the thousands of vendors selling “agentic AI” offer real agents, Gartner estimates



>40%

of agentic AI projects will be cancelled by the end of 2027: rising costs, unclear value, weak risk controls

What gets relabelled: chat assistants, RPA bots and chatbots with a few tool calls.

These are analyst predictions, not measured outcomes.

AI is mainstream — agents are just starting



57%

of German companies with 20+ employees use AI, up from 36% in 2025



11%

of companies engaged with AI already run AI agents; another 29% plan to



66%

of AI users name data protection as a hurdle, the most-cited barrier

EU-wide gap: 17% of small enterprises use AI, versus 55% of large ones (Eurostat, 2025). The barrier is know-how and integration, not models.

Chat + tools + memory: is that an agent?

Often sold as an “agent”

- A system prompt and instructions
- Multi-turn conversation
- Tool calls: search, APIs
- Memory across sessions



What makes it an agent

?

If a human drives every step, it is an assistant.

Is it really an agent? Eight questions



Necessary

All four must be yes, otherwise it is an assistant

1

Goal, not prompt: it gets an objective and a “done” criterion

2

Model controls the loop: it picks the next step and when to stop

3

Acts on systems: it changes state, not just answers

4

Checks itself: it observes results, retries or hands back



Enterprise-grade

Required before anything goes to production

5

Durable state: survives restarts; every step can be replayed

6

Triggers by design: events or schedules, not only a user prompt

7


Bounded autonomy: least privilege, budgets, approvals, kill switch

8

Grounded, evaluated and audited: sources for every claim; drift checks; full transparency; audit log

Five things that get called “agent”

More autonomy → more potential, more risk, more governance needed



Often sold as “agent”

WORKFLOW CASE STUDY

500 submissions to a draft law: a hybrid design

WORKFLOW · STEPS FIXED IN CODE, EVERY STEP TRACED



AGENTIC LAYER · THE MODEL DECIDES



A monitor agent watches for late submissions, decides what to re-analyse, updates the briefing and notifies owners. Changes need human approval (Level 4).

Benchmark to beat (UK Consult): F1 0.76 vs. human coding · 60% of responses needed no edits · median review 23 seconds per response

The different type of agents



Computer-Use & Web Operator Agents

Agents designed to autonomously navigate graphical user interfaces, desktop applications, and web browsers, functioning much like a human user

Good for: Automating workflows that lack dedicated APIs

Watch out: errors over long execution horizons



Personal AI Agents (Cloud-Based)

Intelligent, autonomous systems that maintain a persistent long-term memory across sessions and proactively manage tasks within a user's digital ecosystem.

Good for: Daily organization, managing calendars and emails

Watch out: Deep personalization requires broad permissions



Local / Self-Hosted Agents

Open-source agents running locally on a user's hardware or a private server to ensure complete data sovereignty and isolation from third-party vendor clouds.

Good for: Privacy-conscious users who want to process sensitive data

Watch out: Can pose a major security risk if misconfigured



Enterprise Agents & Specialized Workflows

Agents integrated into corporate systems to manage business processes while adhering to strict governance, role-based access controls (RBAC), and security policies.

Good for: tasks requiring domain-specific expert knowledge

Watch out: coordination overhead and latency

Four patterns behind the word “swarm”



Pipeline

Fixed stages; each step feeds the next.

Good for: document processing at volume

Watch out: this is a workflow, not a swarm



Router / handoff

A triage agent hands the case to a specialist agent.

Good for: service desks, citizen portals

Watch out: context gets lost at each handoff



Orchestrator-worker

A lead agent splits the task; parallel sub-agents research; the lead synthesises.

Good for: broad research across many sources

Watch out: token cost and synthesis errors



Debate / evaluator

Agents critique each other's output; a judge decides.

Good for: red-teaming, risk review

Watch out: gains often small vs. simply sampling more

FIND, ANSWER, ACT

We connect your organization's distributed knowledge into a secure knowledge base. Search, chatbots, and AI agents operate on top of it—always solely with the information each respective user is permitted to see.

YOUR DATA

SharePoint & Intranet

Fileshares

DMS & Electronic Records

Business Applications



Enterprise AI Plattform

The groundwork is done.

Connect data · Inherit permissions · process & index ·
vectorize · Integrate LLMs



YOUR AI-Apps & -Integrations

Find

iFinder – Enterprise Search

Answer

iAssistant – Chat with your data

Act

iHub & AI-Agents

Custom Apps

Your Assistants via API & MCP

Your knowledge. Your rules.

Four things to take home

1

Ask who controls the loop

Chat + tools + memory is an assistant. An agent decides its next step.

2

Start with workflows

Most value today comes from governed workflows with human review.

3

Multi-agent only where it pays

Parallel, high-value work, with a budget for ~15× tokens and stronger evaluation.

4

Build the foundation first

Grounded data, sovereignty, logging and oversight. The model is the cheap part.

#besserFinden

Questions?

Daniel Manzke

daniel.manzke@intrafind.com

