14 October 2026

Michael Gross,
DOT MG

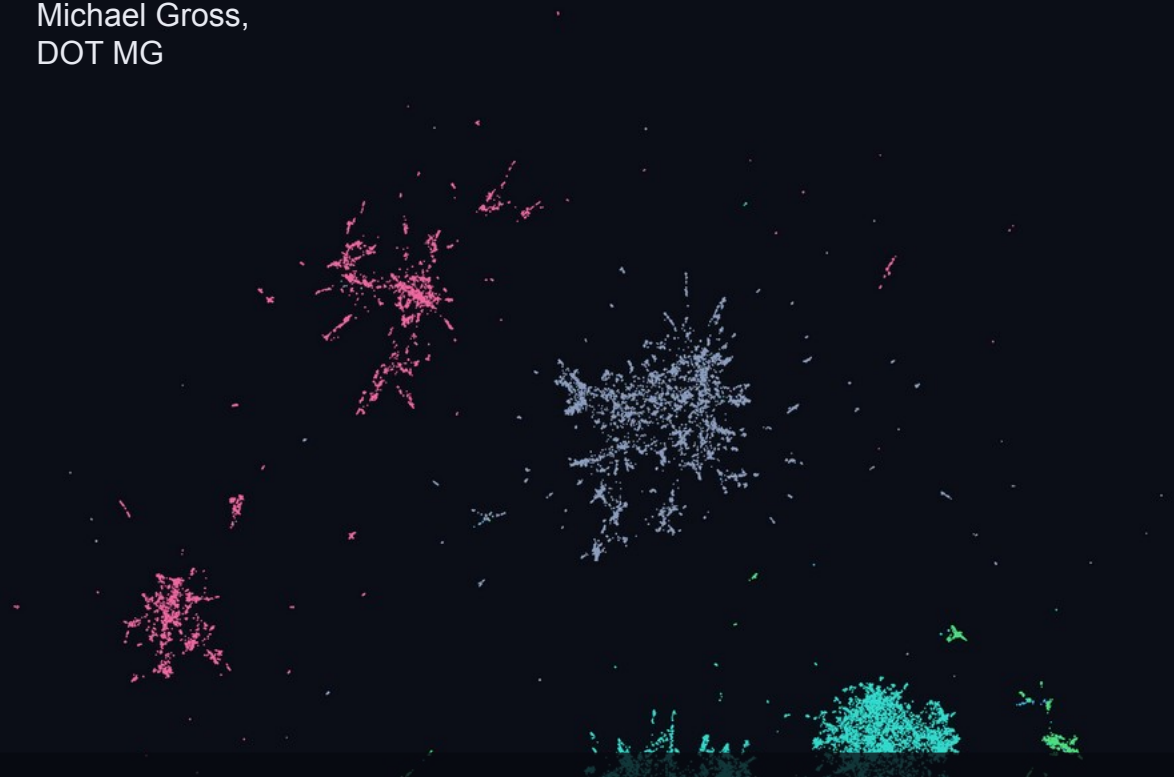

AGENTS NEED HANDS

A personal Second Brain with n8n as its deterministic kernel —
four memories, nine MCP tools, zero cloud

In the Loop 2026 · Berlin

THE ITCH

20 years of data, four silos, zero answers

9,351

scanned documents

Paperless-ngx — invoices, contracts, letters

484

knowledge notes

Obsidian vault — projects, infra, lessons

7,147

watches & jewels

ERP of my antique watch business (PostgreSQL)

1,056

YouTube videos

four channels, fully transcribed with local Whisper

“How much did I pay Telekom in 2024?” — four systems, and not one of them could answer.

LIVE DEMO
the galaxy flies itself · five questions through the cascade

THE DATA MODEL

Four memories, one database

Semantic *what I know*	notes · YouTube · lessons	→ **found by meaning**
Episodic *what happened*	receipts & letters	→ **meaning + exact words**
Business *what a thing is*	a watch's whole life — bought, repaired, sold	→ **exact lookup, SQL**
Meta *how to answer*	72 example questions · every question logged	→ **decides which memory**

One database, three ways to ask — vector, full text, SQL — in a single query.

Markdown is the truth. CrateDB is the compiled index — throw it away, rebuild it.

ROUTING

A deterministic cascade — no LLM decides the route

0 — Known ID

CrateDB node on the ERP mirror — exact

1 — Rule markers

Code node, regex — aggregation intent

2 — KNN majority

embed → CrateDB, 2 of 3 neighbours decide

3 — Hybrid fallback

Switch fallback — never a dead end

72 / 72

test questions, routed correctly

- reproducible
- explainable — the execution log is the audit trail
- cheap — the expensive branch runs last

A new question type = four labeled rows in a table. Not prompt surgery.

WHERE N8N SITS

Every workflow is now an agent tool

3D Galaxy (browser)
webhook — straight in

Agent
Hermes, Claude — via MCP server

→

`frag_second_brain`
THE ROUTER — both channels end here

stages 0–3, no model in the path;
it calls the tools below

→

Data

CrateDB · Paperless API
LM Studio · ERP

…and its eight hands — every one just a workflow, each also callable on its own:

- `wissen_suchen / wissen_lesen` notes, channels
- `belege_suchen / beleg_lesen` receipts
- `semantisch_suchen / offene_posten` semantic · open items
- `uhr_suchen` watches
- `antwort_pruefen` the skeptic

Why this beats a custom API server

- credentials stay in n8n
- edit live, agent stays up
- every call is a replayable execution

LIVE · n8n

the two questions you just saw — the watch and the living room — through their real executions

RELIABILITY

An OS that hides device failure isn't an OS

- **5 services, 5 health checks**

  one try/except hides WHAT broke

- **Outage ≠ “not found”**

  the dead DB used to say “no such watch”

- **A calibrated floor**

  below 0.50 it says “I don't know” — and logs it

Swallowed errors caused every wrong answer this system ever gave me.

TRUST, VERIFIED

A second model, hired to disagree

Shadow mode

runs after the answer — zero latency

gemma audits qwen

different family, different blind spots

Burden of proof

every objection must quote the excerpts

Day one: it caught a receipt from the wrong address. The answering model couldn't see it.

→ live: the workflow itself — and a real verdict from the question log

LIVE · n8n

the skeptic at work — antwort_pruefen and a real verdict on the answer you just saw

THE POINT

Why n8n still matters in the agent era

1

MCP made n8n a toolbox

every workflow you ever built is now callable by a model

2

Determinism where it counts

the reliable half of an AI system is a graph, not a prompt

3

Observability

the log showed my agent bypassing the router

4

Live edits, isolated credentials

fix a tool while Telegram keeps answering

Flowise shut down in August. Their own reason: models outgrew visual low-code for reasoning. Exactly why my LLM never decides the route — the engine isn't there to think. It's there to be the hands.

TAKEAWAYS

Agents need hands. n8n is the hands.

Keep determinism where it counts — spend the LLM only on language.

Make failure visible — an honest “I don't know” beats plausible fiction.

Thank you — questions? Find me in the Community Lounge.

The slides? Yours — for one YouTube sub :)

@itmichaelgross · Michael Gross | DOT MG AI

Questions? Find me in the Community Lounge — right after this session.

Scan → the channel opens with the subscribe button ready. Fair trade — the deck is at michaelgross.net/talk either way.