

Give your team a governed way to make every AI agent follow your standards.

Your teams already use AI coding agents. The open question isn't whether agents can write code - it's whether they follow *your* standards consistently, across every repo. This one day gives your team the kit and know-how to make that happen.

1-day live workshop

Online

Private delivery · scheduled with your team

From €15,000 · excl. VAT

Full refund up to 14 days before the workshop date

01 Why this is worth a day

Right now your AI coding rules live in scattered **CLAUDE.md** files, per-IDE settings, wikis, and personal prompts. They drift, nobody reviews them, and no one can tell you whether agents actually follow them. That shows up as **inconsistent code review, wasted tokens, and governance risk** nobody owns. Writing more rule files doesn't fix it - the missing piece is the operating model around them: versioning, validation, rollout, and measurement.

Not just opinion: models don't use long context uniformly - accuracy drops as input grows, well before the window is full, so much of a bloated context is paid-for noise ([Context Rot, Chroma 2025](#)). Managing that context has become a formally defined discipline - a 2025 survey maps [1,400+ papers](#).

-35%

Token burn per task

-68%

Repeat review comments

31%

Missed-skill activation rate

First design-partner pilot · 12 engineers · 5 repos · 2 weeks. Your own pilot fills in your numbers - the workshop shows you how to measure them.

02 What your team actually brings back

This is not a lecture with slides. It is a build day. Whoever you send leaves with four concrete things:



The full context supply chain, yours to keep

A working kit for governing agent standards: ownership, validation, rollout, recovery, and measurement. Works out of the box.



Trained owners who can run it

Whoever you send operates the kit hands-on - extends it, breaks it, fixes it, ships and rolls back a release - so they can set it up on your stack afterwards.



A two-week pilot path with real metrics

A baseline-first measurement plan: repo selection, adoption evidence, and a leadership report from your own data - not invented figures.



Architecture that outlives our code

Transferable patterns - domain scoping, skill routing, freshness without version bumps - that apply to any agent setup you own.

03 The shift, in plain terms

TODAY, WITHOUT IT
AFTER THE WORKSHOP
AI coding rules
Scattered CLAUDE.md, Cursor rules, wikis, tribal prompts
● One reviewed source, current on every machine
Consistency
Each repo and IDE behaves differently
● Same standards across repos, teams, and tools
Changing a rule
Copy-paste into N repos, hope it sticks
● One reviewed release, staged rollout, instant rollback
Proof it works
"It feels better"
● Pilot metrics you can put in front of leadership

04 A build day, not slideware

59%

of the day is hands-on labs

3h 55m labs · 2h 40m theory
across 8 modules

01 Foundations - why it works

02 The Drift Problem

03 Corpus & Registries

04 Generation & Validation

05 Activation & Routing

06 Distribution & Rollout

07 Governance & the Policy Boundary

08 Pilot, Telemetry & Evidence

05 Who's teaching it



Daniel Glejzner

Principal Consultant & Tech Lead

2× Microsoft MVP · 10+ yrs · 25 clients



Jarosław Żołnowski

Lead Technical Architect

Google Developer Expert · 19+ yrs · 21 clients

06 Who it's for - pick your path

Enterprise-grade mechanics, useful at every scale. Every layer of the context supply chain kit is an independent, dependency-free script - start where you are and add layers as you grow.

SCOPE 01 · ON-RAMP

The owners of your standards

Start with the small group that owns your team's AI standards - your rules, packaged and kept current for the agents you use. Run the lightweight slice first, then add rollout, access policy, and measurement as you grow - without ever migrating to a different system.

Private workshop - core owners

from €15,000

SCOPE 02 · THE PRIMARY PATH

The rollout team

Adoption sticks when the future co-owners train together: the people behind standards, tooling, and governance leave with one shared operating model - not three interpretations of it.

- Staff / principal engineer - owns the standards
- Platform / DevEx engineer - owns rollout and installs
- Engineering manager / security - owns governance and sign-off

Private workshop - the rollout team, one room

sized on the call

SCOPE 03 · GO COMPANY-WIDE

Train everyone, not just the owners

A handful of people own the platform - but knowing how agents use context is fast becoming core engineering fluency, and every developer who gets it ships better with AI. A private workshop trains your owners and builders together and includes a 2-week implementation pilot to prove the rollout in your repos.

Private workshop - company-wide

sized on the call

Included for every participant: 30-day follow-up Q&A call with the instructors

Scope and date: set on a 30-minute planning call

Prices excl. VAT · invoicing, PO, and vendor onboarding supported

Recommended next step

Book the planning call and pick the scope that fits - the core owners, the team who'll run the rollout, or your whole engineering org. Whichever you choose, ask for the same deliverable: **evidence from your own repos**- not enthusiasm. If it holds up, you'll have the evidence - plus the kit and know-how - to roll it out for real.

Book a planning call →

<https://aicontextengineering.eu/> · Private company workshops only - individual seats are closed and the next public cohort will be announced separately. Prices excl. VAT · final scope and price are set on the planning call. This brief avoids invented ROI figures on purpose: the numbers that matter come from your own pilot.