

# Context pilot for teams already using AI.

We help software teams turn scattered repositories, tickets, docs, incidents, and chat decisions into a **source-backed context layer** that AI tools can use without bypassing review, permissions, or production discipline.

## Good fit when

- engineers already use AI coding tools, but answers depend on stale docs or tribal memory
- agents need repository, ticket, incident, or runbook context without broad access
- leadership wants speed, but not silent production writes or unverifiable decisions

01

## Audit one workflow

Pick one real loop: release risk, support-to-fix, incident follow-up, PR review, or agent-assisted change.

02

## Ship one useful context loop

Small implementation or prototype, with boundaries clear enough for a continue/stop decision.

## What we inspect

Repos, ticket flows, docs, incident notes, runbooks, ownership, deploy rules, and the places where AI is already being used informally.

No need to open the whole company. A narrow workflow is usually better.

## What we protect

Source trails, freshness, permission boundaries, human approval for risky writes, and a visible diff between old state and proposed state.

The point is not more automation. The point is safer speed.

## What you get

A workflow map, risk boundaries, one practical context loop, and a recommendation on what to automate, assist, or leave manual.

Enough evidence for the next decision.

## PROOF POINT

### PLAYGRND

PUBLIC RECORD + AI-ASSISTED BUILD

PLAYGRND is our public sports-record product. It combines fast SSR pages, a Go API, Postgres/Redis, WhatsApp magic links, claim/correction flows, and derived season aggregates. AI helps the team move faster, but official records still require clear sources, controlled writes, fallback paths, and human confirmation. That is the same operating model we bring to client AI workflows.

## Contact

hello@hills-lab.hr  
hills-lab.hr

## Company

HILLS LAB d.o.o. · Split, Croatia  
VAT ID: 48104693060