CONCEPT BROCHURE, NOVEMBER 2026

CONCEPT



iogentic

Agent Platform

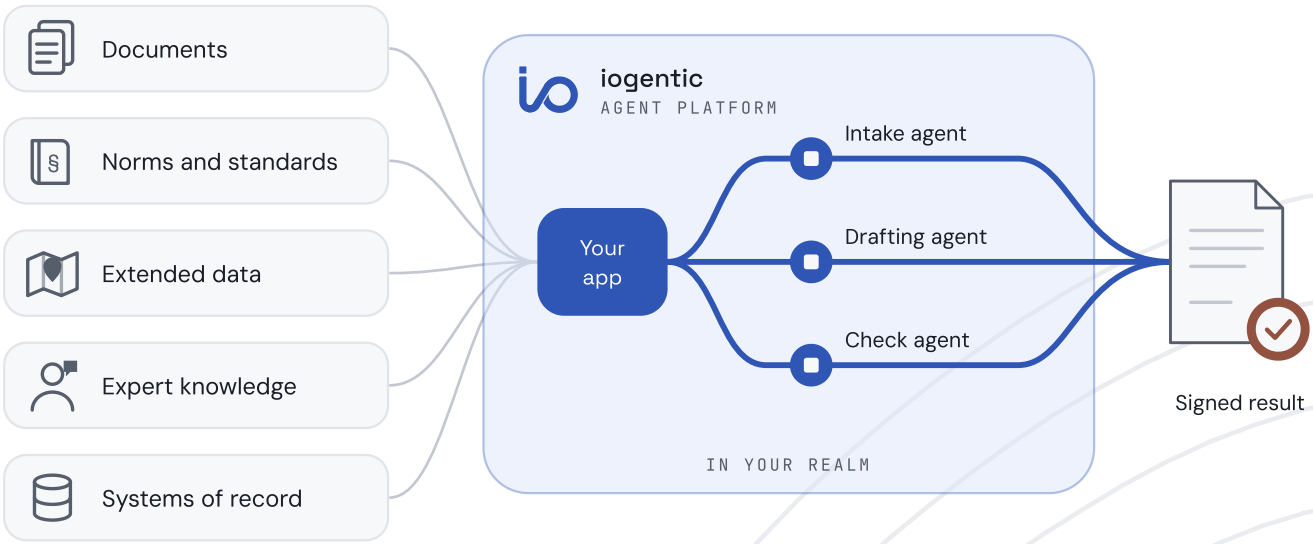
AI agents for planning and infrastructure work. We plan to build each agent with the people who do the job today, and a person signs each result before anyone uses it.

This is a concept. Nothing here is released yet, and the worked example uses made-up data. The photos are generated images.

Purpose

iogentic is designed to run AI agents for one specific job in your organisation, for example putting together the tender package for a cable route. In the design, the agents read your documents and systems, write a first draft and check it. Your people decide the rules at the start and sign the result at the end.





Workshop first

We start on site with the people who do the job today. Together we write down the steps, the terms and what a good result looks like, before an agent drafts anything.

Your realm

In the design, every organisation gets its own closed space for data, models and agents, with no route to another realm.

Checks before drafts

We plan to test every step with code checks and with errors we plant on purpose. When a case fails, it becomes a new test once a person has approved it.

Two signatures

In the design, one person signs off the input and a second person signs off the release.

STATUS, 24 NOVEMBER 2026 The platform is a concept. The harness is in development and the runtime has no code yet. We are five founders in Germany and have spoken with one operator so far.



02 APPROACH

Six Steps

In the plan, it starts with a workshop on site, where we agree with your team on what will be checked.

1 2 3 4 6

Workshop

We sit down with the tender author and the project lead and go through the steps, the terms, the norms and the cases that usually go wrong. Every open question gets an owner.

Map

For each step we note which system it reads, which data it needs and what a good result looks like. The map also sets what each agent may read.

App

The workflow becomes an app in your realm, with one agent for intake, one for drafting and one for checking.

Validate

Before a person sees a draft, code checks and planted errors are meant to run against it.

Sign

Two people sign: the tender author at Gate 1 and a project lead at Gate 2.

Tune

When a case fails, it becomes a new test, but only after a person has approved it.



03

WORKED EXAMPLE

Tender Package

A tender package for a 20 kV cable route, from the first workshop to the signed release.

The project is synthetic and the run is scripted. No model was called and no customer data was used.

SYNTHETIC, SCRIPTED

Nine Phases

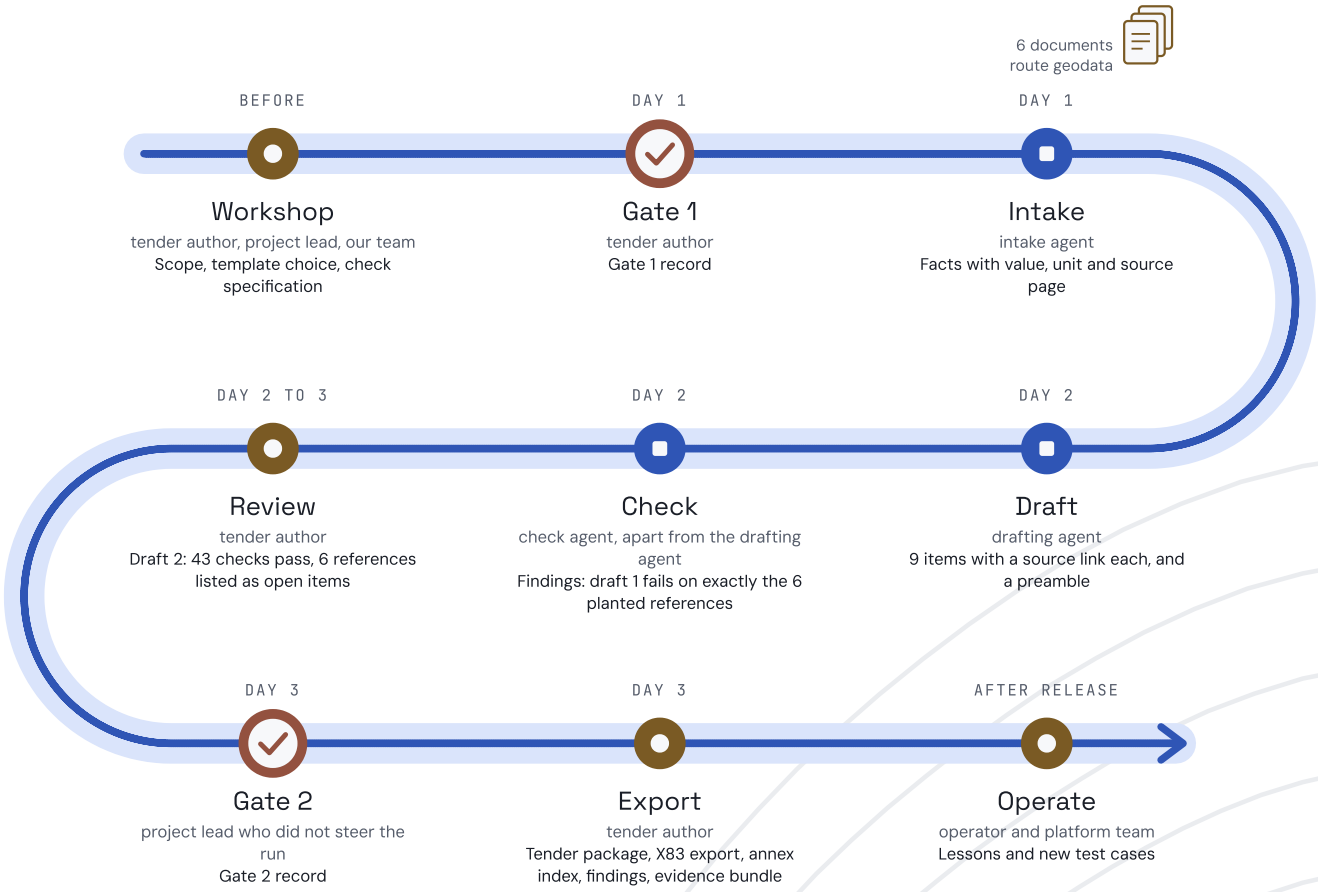
2,663.8 m
cable route, 20 kV

17
crossings: paths, a main road, a railway and waters

6
project documents with 14 pages

9 of 31
template items apply

person · agent · gate with a signature · Days are an example



Drafts go into a separate draft area, so your own systems stay the system of record. The days come from the scripted run and nothing was measured. The findings in the run are scripted too, and a domain expert has not reviewed the gold set yet.



04 GATES

CONCEPT, NOT YET RUN

Two Gates

In the design, the run stops at each gate and waits for a signature. The tender author signs the scope, the inputs and the check specification. A project lead who did not steer the run signs the release.



SIX RECORDS PER RESULT, PROPOSED

- 01 Gate 1 record
- 02 Task record
- 03 Trace and prompts
- 04 Models and tools
- 05 Check results
- 06 Signed release

Who signed

The Gate 1 record names who signed, in which role and when, plus the hash of what they signed.

Who checked

In the proposal the check agent works separately from the drafting agent, and the tender author goes through its findings in the review.

What is kept

We propose six records for every result. Where they are stored is something to agree with your data protection officer.

CONCEPT, NOT BUILT

Agentic OS

The design has three layers. The OS gives every workflow the same eight services. A realm is one organisation's closed space. An app is your workflow with its own screen, and apps only talk to each other through typed events on the bus.



Scheduler

A run is a process with nine states, from created to done.

Memory

Knowledge is read from your systems. Memory is written by agents and approved by a person.

Shelves

Agents and capabilities install like packages, with a manifest.

Model gateway

Allowed models and regions per realm, model and version logged per call.

Connectors

Drivers to your systems, with allowed tools per connector.

Audit log

Append only, kept apart from the evidence.

Event bus

Typed events: step, gate waiting, gate signed, evidence added, export.

Test bench

Code checks and model judges run against planted errors.

MEASURED, SYNTHETIC DATA

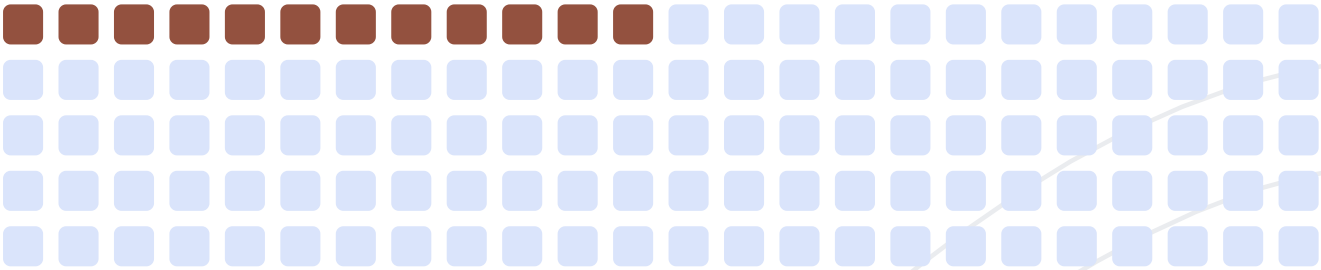
Evaluation Run

We tested five current models on two synthetic tender projects on 8 November 2026, each inside an agent loop with tools. None of the five models passed every task we could grade. In this test, none of them got the work right on its own.

The design answers this with expert knowledge from your team, checks between the steps and your rules as model context. We have not yet measured how much each helps. Code checks look for numbers that no tool returned, wrong units and broken references. The bar is 80 % of attempts at grade 4 of 6 or better, with a mean grade of 4 or better.

120

attempts in one run: 5 models, 4 tasks, 2 projects and 3 repeats. One square per attempt.



12 of 120 attempts were capped at grade 3, because the answer held a number that no tool had returned. The squares are grouped, so their position means nothing.

12

code checks run on every attempt, among them schema, structure, units, quantities within 5 % of the reference, references, cited value on the cited page, allowed norms.

0.5

The two model judges reached a kappa of 0.5, a measure of agreement corrected for chance. Our bar is 0.6, so for now no judge grade counts.



Our evaluation run of 8 November 2026 on two synthetic tender projects, results counted on 24 November 2026. The projects are synthetic, so the numbers apply to this test only. The task verdicts come from the evaluation report of that run. The gates are still a concept.

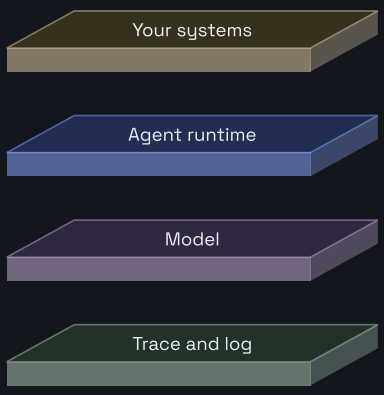


07 SOVEREIGNTY

DESIGN GOAL

Where It Runs

Data, model and log can stay in the EU. You choose where each layer runs.



Your systems	Stay where they are. Agents read through connectors, read only in the example.
Agent runtime	A data center in Germany, an EU cloud region, or your own servers.
Model	Open weights you host yourself, an EU-hosted service, or an API in an EU region by opt-in.
Trace and log	Exported as files, and you set the retention.

What stays with you

Plan sets, templates, tender packages, check specifications and domain knowledge. In the design, no data goes to an external model service unless you opt in.

Limits

We are not lawyers and hold no certificate. Whether you comply depends on the use case. In the design, the records are meant to help you show it.

Four Steps

You can stop after any of the four steps. The first pilot would cover one tender package, from complete plan sets to a signed draft, running next to your normal process. Steps three and four need the runtime, which is a concept today.

STOP AFTER ANY STEP

Bottleneck check
A FEW DAYS

A map of where agents could pay off in your work, and a look at the data.

Trial run
ONE TO TWO WEEKS

A first draft on one real case, partly on site, with an estimate and how we got to it.

Pilot
A FEW MONTHS

Released results, signed by named people in your organisation.

Operation
ONGOING

In your environment or a German sovereign cloud, run by us or by your team.



What you keep

- A map of the work
- A draft on a real case
- Signed results with six records each
- Documentation of how each agent is built and tested
- Your plan sets, templates, check specifications and domain knowledge

What the pilot measures

- Author days per tender, from a time log
- Errors in annex references
- Share accepted at Gate 2 in the first round
- Cost per tender: model, hosting and author hours

We set prices, durations and stop thresholds together with you before the first round.



iogentic

Talk to Us

Thirty minutes, no slides. Tell us about one job in your organisation, and we will look at it together to see whether an agent could help.

hello@funkeltronic.ai

Concept brochure, 24 November 2026. The platform is not released. All photos are generated images. The worked example and the evaluation run use synthetic projects. The numbers come from our evaluation report of 8 November 2026 and the worked example on iogentic.ai.