WHITEPAPER, NOVEMBER 2026

CONCEPT



iogentic

Agentic Workflows

Regulated industry, with construction tenders as the example

How AI agents can do document work in German construction and utilities, and how their results are checked before a person signs them. The worked example is the bill of quantities of a public tender.

This concept paper comes from the founding team of iogentic. The platform is not released. The worked example and our evaluation use synthetic projects, and the photos are generated images.

In Brief

The main points, each with its source.

German companies are adopting AI, and a growing share of them are looking at agents: software that reads documents, calls tools and drafts work. Tenders, permits and safety records follow written rules and carry the signature of a named person, so a draft has to be checked before anyone relies on it. This paper looks at one case, the bill of quantities of a public construction tender, and describes how the planned iogentic agent platform addresses it.

01 **Adoption**

37 percent of German companies with ten or more employees used AI in 2026 [1]. In a Bitkom survey, 11 percent of the firms that use or plan AI already use AI agents [2].

02 **Tenders**

Almost half of the 199,334 public awards reported for 2024 were construction contracts [3]. Each of them has to describe the work so that every bidder reads it the same way [4].

03 **Model results**

In our run, none of five current models passed every task we could grade [5]. In published benchmarks, agents that solve a task once often fail when the same task is repeated [6, 7].

04 **Workflow design**

The design fixes the order of work, writes the checks before drafting, gives every drafting agent a separate checker and puts a person at each gate [8, 9].

05 **Records**

The AI Act asks high-risk systems for logs, technical documentation and human oversight [10, 11, 12]. The platform is planned to keep six records for every result, so a reviewer can follow what happened.

CONTENTS

This Paper

01 Regulated Work 04

German companies are adopting AI, while their document work follows written rules and ends with a signature.

02 Public Tenders 06

Public construction is expected to grow faster than the sector as a whole, and its tenders describe the work in a bill of quantities.

03 Model Limits 09

Our evaluation run and published studies.

04 Agentic Workflows 11

Workflows and agents in one run.

05 The Platform 12

A planned runtime for agents, with shared services in an OS, one closed space for each organisation and one app for each workflow.

06 Harness 15

Design study of the interface.

07 Worked Example 18

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

08 Sovereignty 22

Where data, model and log run.

09 Getting Started 23

A pilot in four steps.



01

CHAPTER 01

Regulated Work

German companies are adopting AI, while their document work follows written rules and ends with a signature.

Adoption and Rules

AI use, staff and the rules in force in November 2026.

37 %
of German companies with ten or more employees used AI in 2026 [1]

11 %
of firms that use or plan AI already use AI agents, Bitkom survey [2]

25.4 %
of employees in shortage occupations are 55 or older [13]

Adoption

AI use in German companies has grown. Destatis counts 37 percent of companies with ten or more employees using AI in 2026, and 73 percent of companies with 250 or more employees [1]. Bitkom reports 57 percent for firms with 20 or more employees [2], and a news report on an ifo survey gives 54.5 percent in June 2026 [14]. The surveys use different bases, so the figures differ.

Agent use is still low. Of the Bitkom respondents that use or plan AI, 11 percent use AI agents and 29 percent plan to [2]. In construction, AI use started later than in other sectors. According to the same news report, AI use in the main construction trade rose from 7.1 to 39.8 percent within three years [14], while generative AI was still below 15 percent in construction in 2025 [15].

Staff

Skilled staff for tenders and plans are in short supply. In the relevant construction occupations, 70,323 skilled workers were missing on average in 2023 [16]. In two core occupations of building planning, 27.9 and 25.3 percent of the employees are 55 or older [13]. Labour productivity in German construction has hardly grown since the 1990s [16], and staff in architecture and engineering offices spend about a quarter of their time on work outside their projects [17].

Rules

Several rules for AI changed in the same period. The AI Act applies generally from 2 August 2026, with later dates for some high-risk systems [18], the Data Act has applied since 12 September 2025 [19], and the German NIS2 act was promulgated in December 2025 [20]. Inside a company, the works council has a say when technical systems can monitor staff, and it has to be informed about work procedures that use AI [21, 22].



02

CHAPTER 02

Public Tenders

Public construction is expected to grow faster than the sector as a whole, and its tenders describe the work in a bill of quantities.

Market and Tenders

Construction volume, public awards and the rules for tender documents.

544.48 bn €
construction volume in Germany 2025, DIW estimate [23]

almost 48 %
of 199,334 public awards in 2024 were construction contracts [3]

10 + 4 %
of the planner's fee for buildings: preparing and assisting the award [24]

Market

The DIW estimates the German construction volume for 2025 at 544.48 billion euro, of which 97.62 billion euro are public construction [23]. For 2026 and 2027 it expects public construction to grow by 6.7 percent a year in real terms, faster than the sector as a whole [23]. The construction trades counted 20,857 establishments with 20 or more employees in 2024 [25].

Tenders

German contracting authorities reported 199,334 awards worth 135.2 billion euro for 2024. Construction contracts made up almost 48 percent of the count and 47.7 billion euro of the value [3]. Above the EU threshold, 5,404,000 euro for works since January 2026, the European rules apply [26, 27].

VOB/A requires the works to be described clearly and so completely that all companies understand the description in the same sense [4], and the Act against Restraints of Competition (GWB) asks for a description that is as unambiguous as possible [28]. Under the HOAI, preparing and assisting the award carry 10 and 4 percent of the planner's fee for buildings [24].

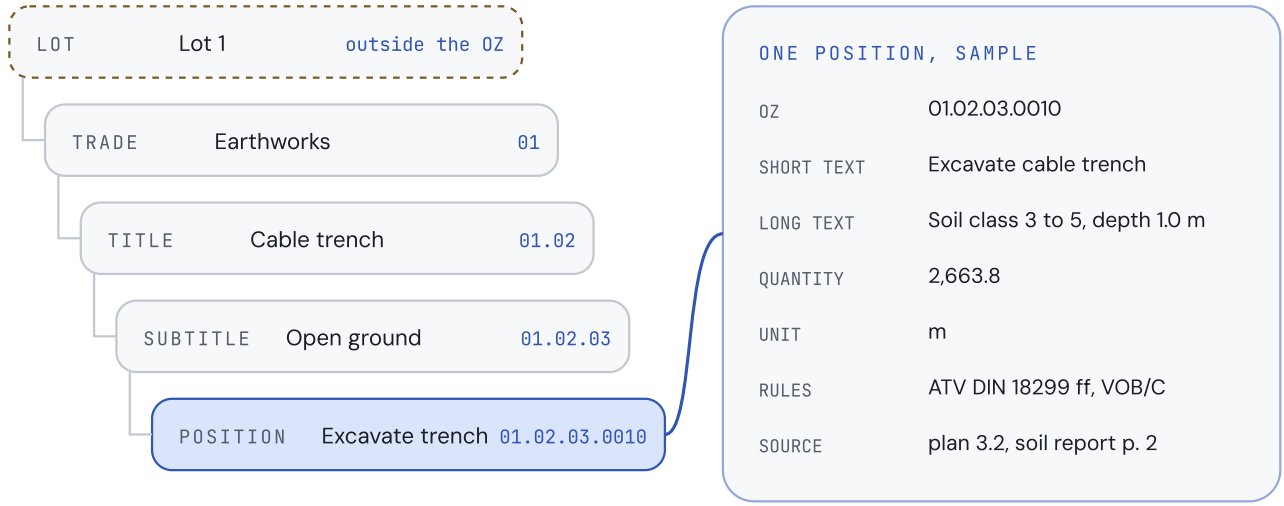
Capacity

Public bodies posted almost 12,400 jobs in construction occupations at all qualification levels in 2025, 73 percent of them in municipalities [29]. A 2017 review of five completed building projects in Bremen named errors in the bill of quantities as one of five groups of reasons for change orders [30]. It suggested a consistent four-eyes review of the bill, which needs time and skills [30].

SAMPLE POSITION

Bill of Quantities

Levels, positions, standard texts and the GAEB exchange phases.



A bill of quantities (Leistungsverzeichnis, LV) is a hierarchy. Trades, titles and subtitles group the work, and the positions at the bottom describe it, each with a quantity, a unit and a reference to the technical rules of VOB/C [31]. The Ordnungszahl numbers every element, and GAEB DA XML allows up to five levels and 14 digits [32]. A tender can also be split into lots, which one operator's guideline keeps outside the Ordnungszahl [33].

Standards define the texts. STLB-Bau provides standard texts for positions and is compiled by the working groups of GAEB [34]. ATV DIN 18299 sets the general rules for all construction work [35], DIN 276 structures the costs [36], and DIN 18290 links building models to GAEB bills of quantities [37].

GAEB DA XML is the exchange format, currently in version 3.3 with a beta of 3.4 [38]. Its phases describe who sends what: X83 is the request for offer from client to bidder, X84 the offer that comes back and X85 a side offer [39]. One state building authority, for example, does not let bidders change the structure of the file [40] and still required version 3.2 in 2024 [40].



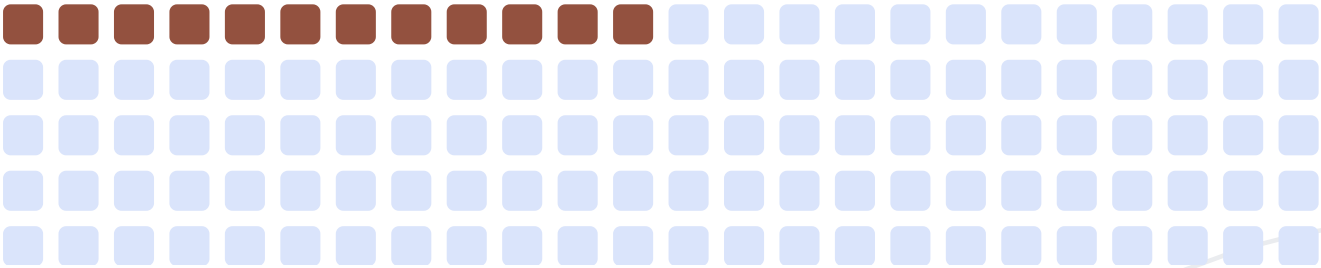
MEASURED, SYNTHETIC DATA

Model Limits

Our evaluation run of 8 November 2026 on synthetic projects.

We tested five current open-weight models on two synthetic tender projects in November 2026. Each model worked inside an agent loop with tools, on four tasks, three times per project, so 120 attempts in all [5]. The tasks are the ones a tender author does: extract project facts, check references, draft positions of the bill of quantities and write the preamble.

No model passed every task we could grade. One model passed the project facts and another the positions. The references were the easiest task, and four models passed it. In 12 of the 120 attempts, the answer held a number that no tool had returned [5].



One square per attempt. Red: 12 attempts with a number that no tool returned, grouped; their position means nothing.

MODEL	T1 PROJECT FACTS	T2 REFERENCES	T3 BILL OF QUANTITIES ITEMS	T4 PREAMBLE
Model A	3.5 fail	3.2 fail	4.5 fail	5.2 no verdict
Model B	3.9 fail	6.0 pass	5.4 pass	5.5 no verdict
Model C	2.9 fail	5.5 pass	3.8 fail	5.3 no verdict
Model D	4.2 pass	6.0 pass	4.4 fail	6.0 no verdict
Model E	3.9 fail	6.0 pass	4.9 fail	6.0 no verdict

Mean grade of six attempts, scale 1 to 6. Pass: mean 4 or better and 80 % of attempts at 4 or better. T4 has no verdict because the model judges are not calibrated. Model names withheld. [5]

Published Work

Studies on agent reliability, invented answers and model judges.

Published studies also report reliability problems. On tau-bench in 2024, a leading agent solved fewer than half of its tasks and got less consistent when the same task was repeated [6]. A 2026 study of 15 models found that recent gains in capability brought only small gains in reliability [7]. When models answer questions about documents they were given, the share of invented answers grows with the length of the context [41].

A second model as checker has limits. A strong model judge can agree with people about as often as people agree with each other [42], but plain agreement rates overstate this, because they ignore agreement by chance [43]. Our own two judges reached a kappa of 0.5, below the bar of 0.6 we set, so no judge grade counts yet [5].

Our approach rests on the three measures below. We have not yet measured their effect in a run. The right context alone already makes answers more factual [44], and better retrieval cut failed look-ups by two thirds in one published test [45].

Expert knowledge

The tender author's rules, terms and failing cases, written down in the workshop before any agent drafts.

Workflows with checks

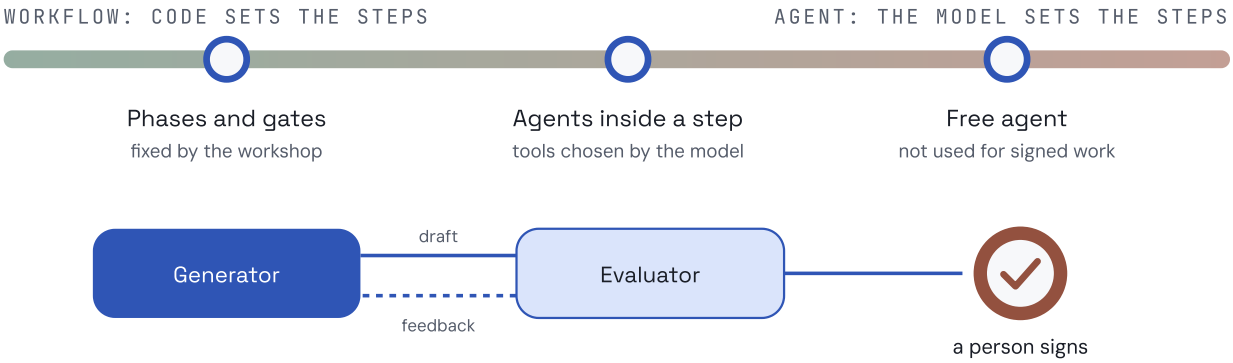
Fixed phases, a check agent apart from the drafter, code checks on every result and people at the gates.

Models with context

Your rules, norms and past cases go to the model as context, read through connectors.

Agentic Workflows

Workflows and agents in one run.



Anthropic distinguishes two kinds of systems. In a workflow, code sets the path and the model fills in the steps. In an agent, the model decides its own next step and which tool to use [8]. Regulated work needs both. The order of work is fixed by rules, but reading a soil report or matching a plan to a template item inside a step needs judgement.

The design draws on three published patterns. In the evaluator-optimizer pattern, one model call drafts and another checks and sends feedback [8]. Human checkpoints pause the run where a person has to decide [8], and OpenAI describes human intervention as a safeguard for high-risk actions [9]. Evaluations set a baseline before anyone tunes a prompt [9], and they should judge the path of tool calls as well as the final answer [46].

Agents use tools through a standard interface. Anthropic introduced the Model Context Protocol in November 2024 as an open standard that connects models to data and tools [47]. Its specification requires the user's consent before a tool is used [48].

ELEMENT	IN THE DESIGN OF THE PLATFORM
Fixed phases and gates	Set in the workshop. A run cannot skip a gate.
Drafter and checker	The check agent works apart from the drafting agent and sends findings back.
People at gates	Gate 1 before drafting, Gate 2 before release.
Evals before tuning	Planted errors and code checks run before Gate 2.
Tools through MCP	Connectors to your systems, read only in the example.



05

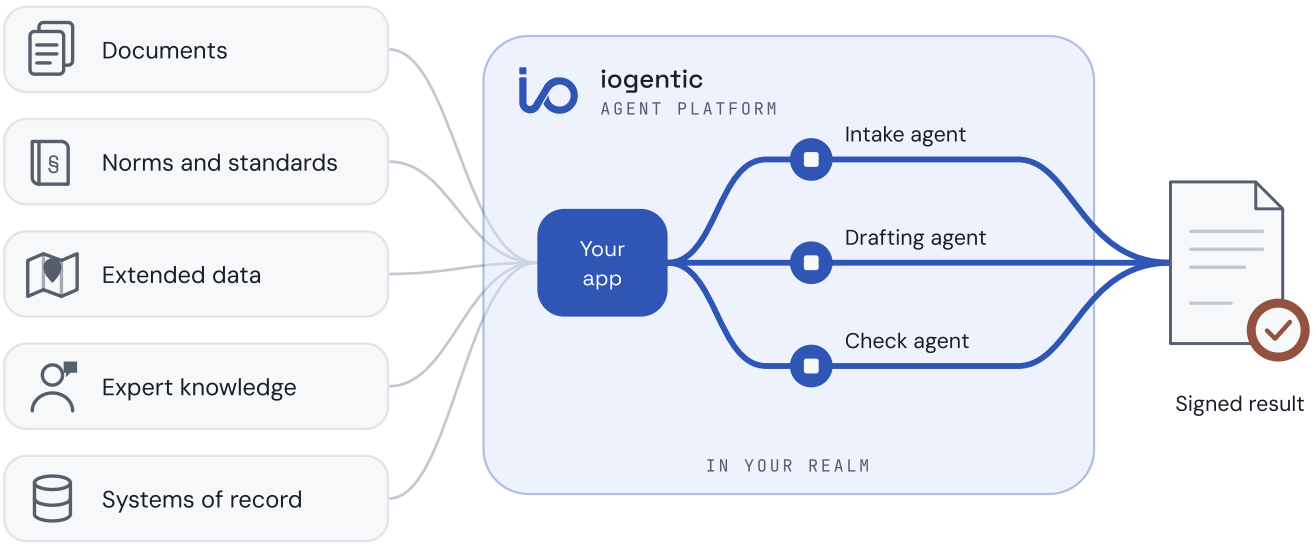
CHAPTER 05

The Platform

A planned runtime for agents, with shared services in an OS, one closed space for each organisation and one app for each workflow.

Architecture

From the inputs to the signed result.



In the plan, we first learn the workflow in a workshop with the people who do the work. We then build it as an app on an agentic OS and test each step before anyone relies on it. Where a rule asks for a person, the run is designed to wait for a signature.

The platform consists of three products. The iogentic harness is the app in which people build, run and check agents, and it is in development as a native app. Pipesmith (pipesmith.ai) is meant to build, verify and store agents; its CI and registry are in development and its QA is a proposal. The iogentic runtime is planned to start an agent only from a verified build, run it in a sandbox and record every step. It is a concept without code.

PRODUCT	STAGE	STATUS ON 24 NOVEMBER 2026
iogentic harness	Develop	In development, native app
Pipesmith	Build, verify, store	CI and registry in development, QA a proposal
iogentic runtime	Run	Concept, no code yet

CONCEPT, NOT BUILT

Agentic OS

In the design, the OS provides the same eight services to every workflow.



A realm is the closed space of one organisation, with its own data, secrets, models and agents, and with no route to another realm. An app is one workflow with its own screen. Apps exchange data only through typed events on the bus, and your systems stay the system of record.

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

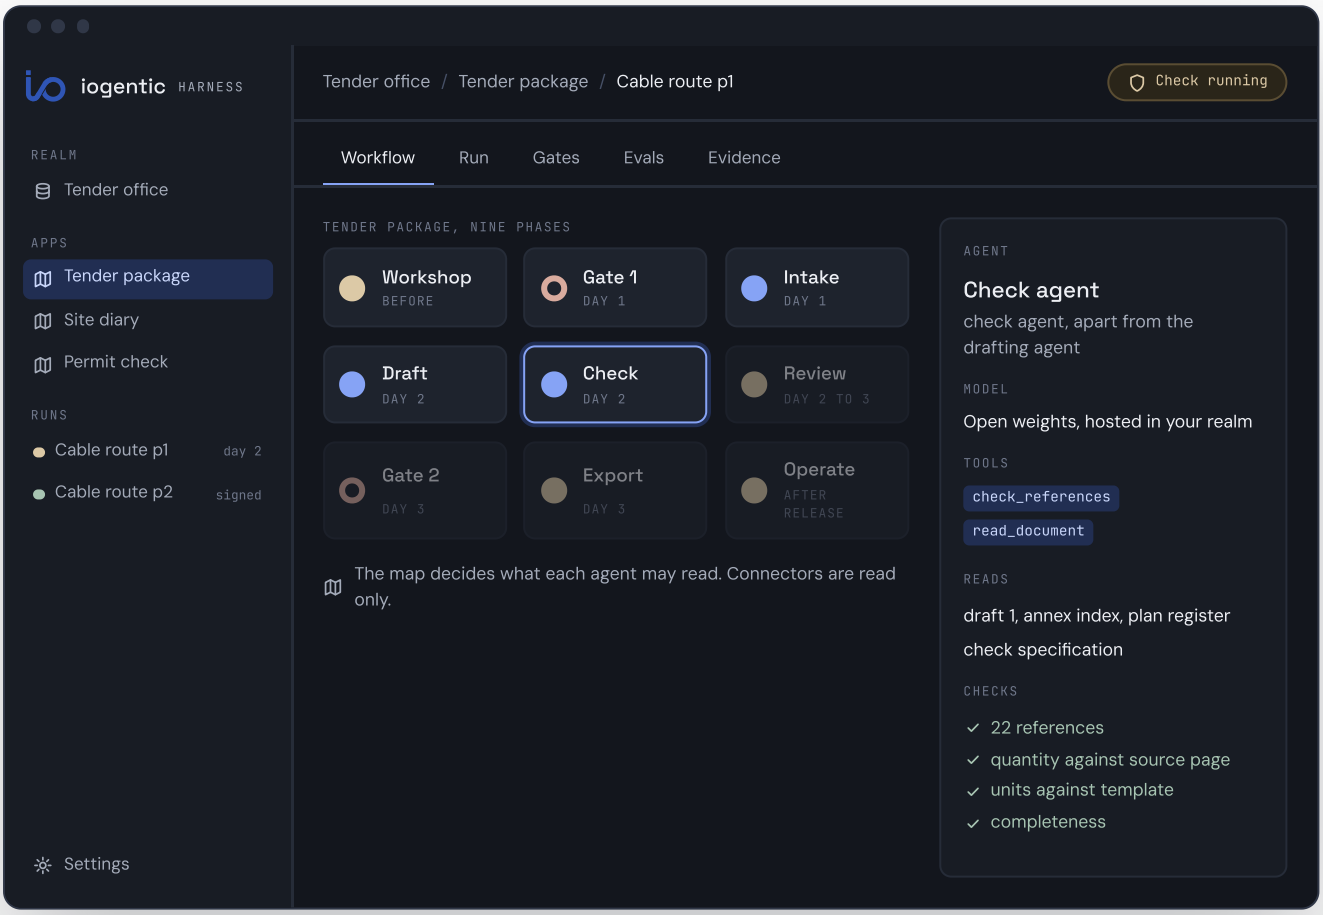
Code checks and model judges run against planted errors.

UI STUDY, SAMPLE DATA

The Harness

Design study of the harness as the interface to the platform.

The harness is the interface through which people use the platform. It is a native desktop app in development and shows the conversation with an agent and every tool call with its timing, and it asks for approval before the agent changes anything. The design study below uses sample data and shows four screens: the workflow of an app, a run, the signature at Gate 2 and the evaluation matrix. The study is not built.

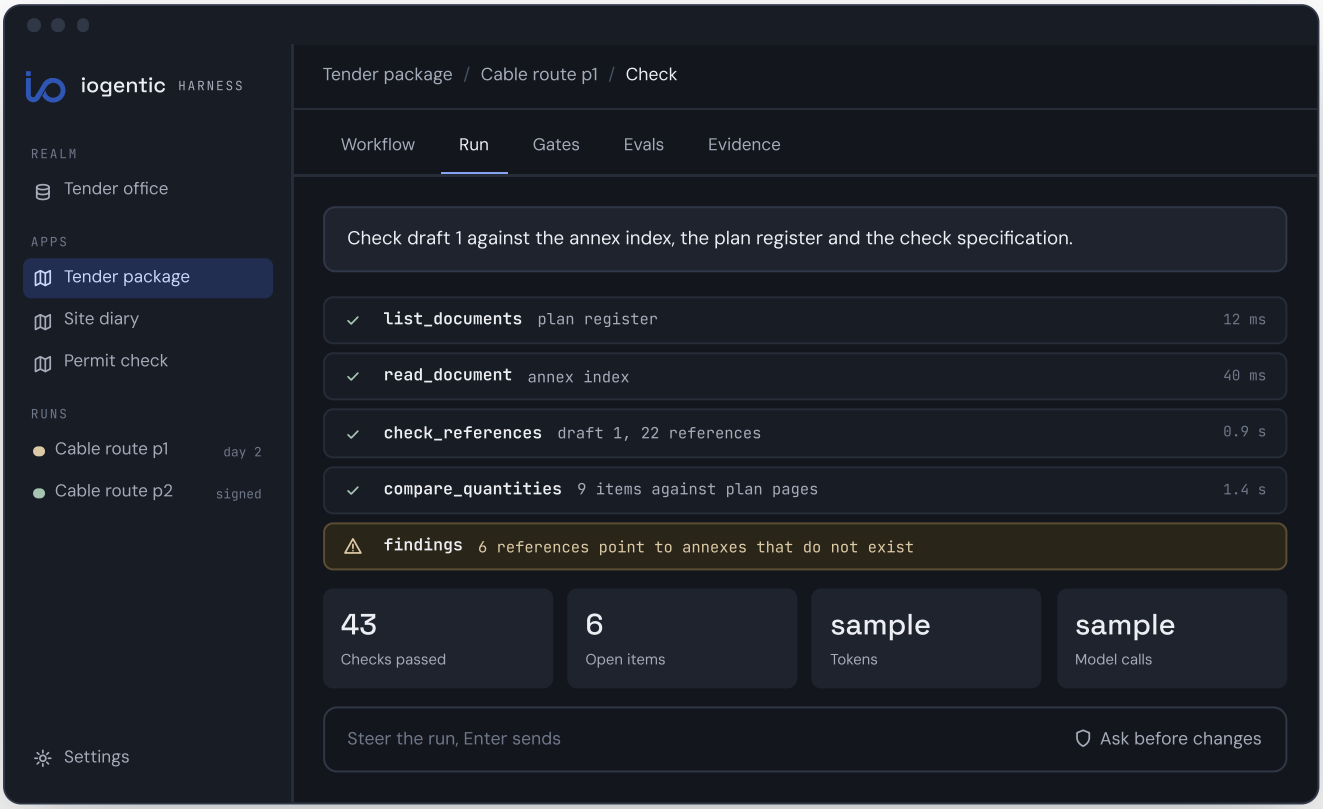


Workflow The inspector lists the tools, reads and checks of the check agent.

UI STUDY, SAMPLE DATA

Run

A run of the check agent on the first draft.

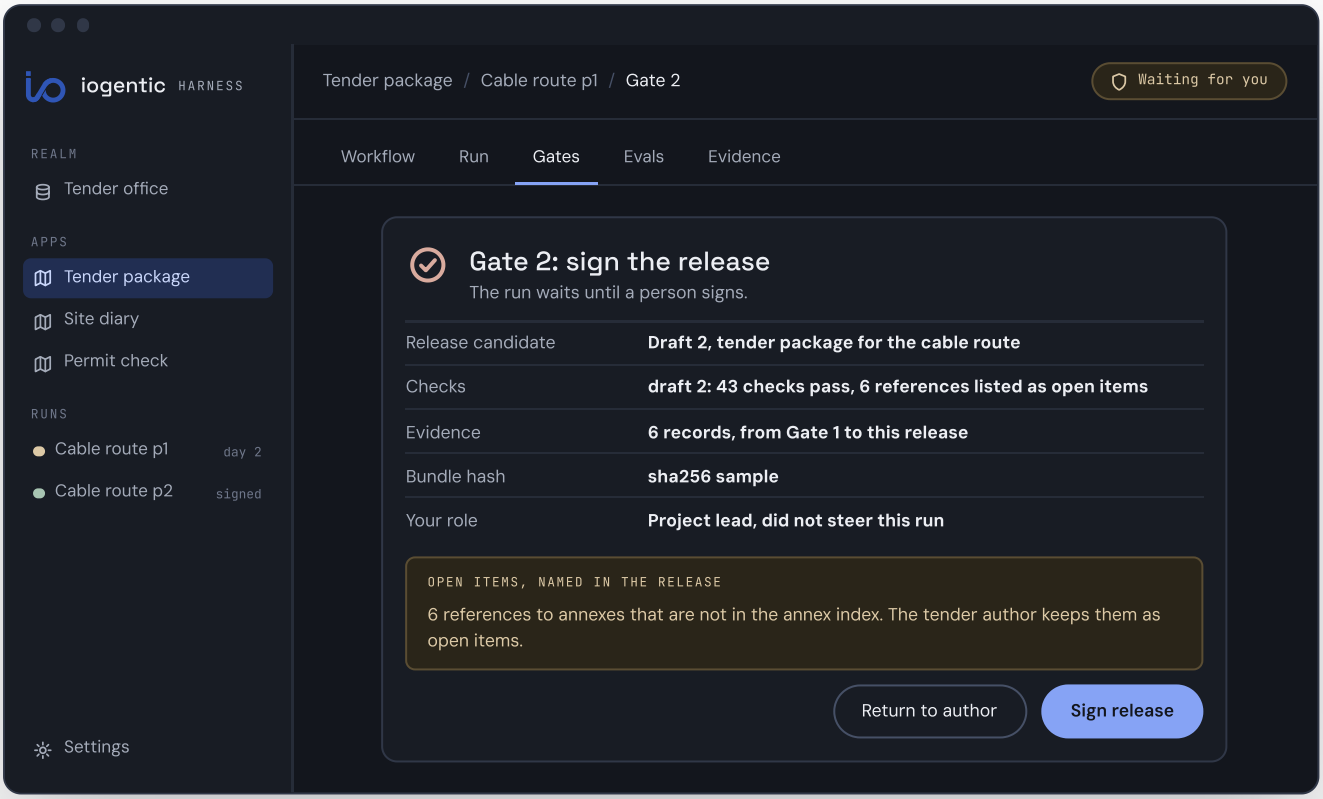


Run Every tool call with its timing. The findings of the check agent go to the tender author.

In the design, every run shows the task, each tool call with its timing, the checks that passed and the findings. Where a rule asks for a person, the run waits and shows that it waits for a person. The team can steer a run while it works, and any change to files needs an approval first.

Gate and Evaluation

The signature at Gate 2 and the evaluation matrix.



Gate 2 Signature with the open items named.



Evals The matrix of our run of 8 November 2026, measured on synthetic projects, with anonymised models. [5]



07

CHAPTER 07

Worked Example

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

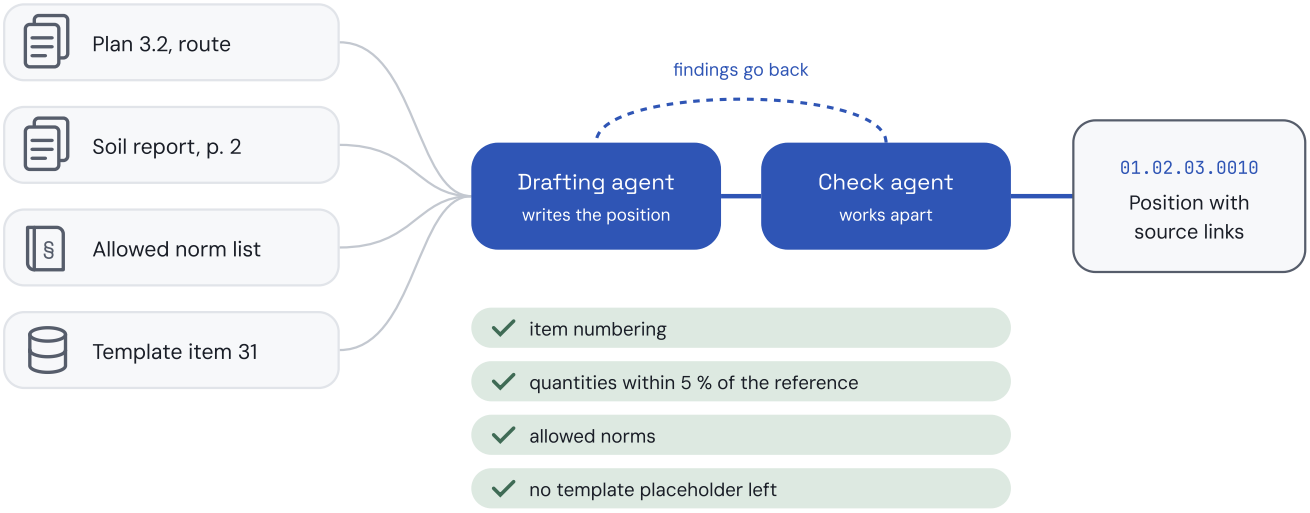
Phases, people and gates of the scripted run.



20 kV cable route, 2,663.8 m, 17 crossings: 11 paths, 1 main road, 1 railway, 4 waters. 25 facts, 22 references with 6 planted broken ones, 9 items. The gold set is not yet reviewed by a domain expert. [49]

Drafting a Position

How one position is drafted and checked.



In the example, every position starts from sources: the plan, the soil report, the list of allowed norms and the template item. The drafting agent writes the position with a link to each source page. The check agent works apart from it, compares quantities with the plan pages and units with the template, and sends findings back [8]. Positions that fail go back to the drafter, and the tender author sees the positions that pass together with the findings.

Records

The six planned records of a result and the AI Act.



In the design, the run keeps six records for every result, from the Gate 1 record to the signed release. They show who set the scope and who checked the result.

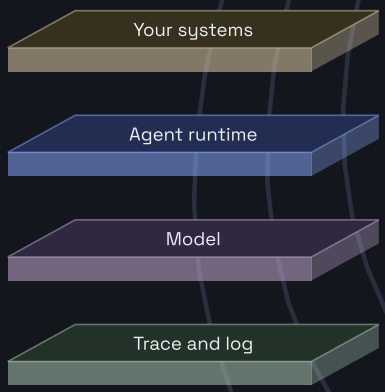
The AI Act asks high-risk systems for a risk management system, technical documentation, automatic logs and effective human oversight [50, 10, 11, 12]. Whether an agent that drafts tender documents is high-risk has to be assessed case by case. For high-risk systems listed in Annex III, the obligations apply from 2 December 2027, after the Digital Omnibus moved the date [18, 51]. The table maps four of the six planned records to these articles. This is our own reading, without legal review.

RECORD	WHAT IT HOLDS	AI ACT
Gate 1 record	who set the scope, inputs and check specification	Art. 14, human oversight [12]
Trace and prompts	every step, tool call and model call	Art. 12, record keeping [11]
Models and tools	which model and which tools, with versions	Art. 11, technical documentation [10]
Check results	code checks and planted errors against the requirement	Art. 9, risk management [50]

DESIGN GOAL

Sovereignty

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



Your systems	Stay where they are. Agents read through connectors.
Agent runtime	Data center in Germany, EU cloud region or your own servers.
Model	Open weights you host, 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.

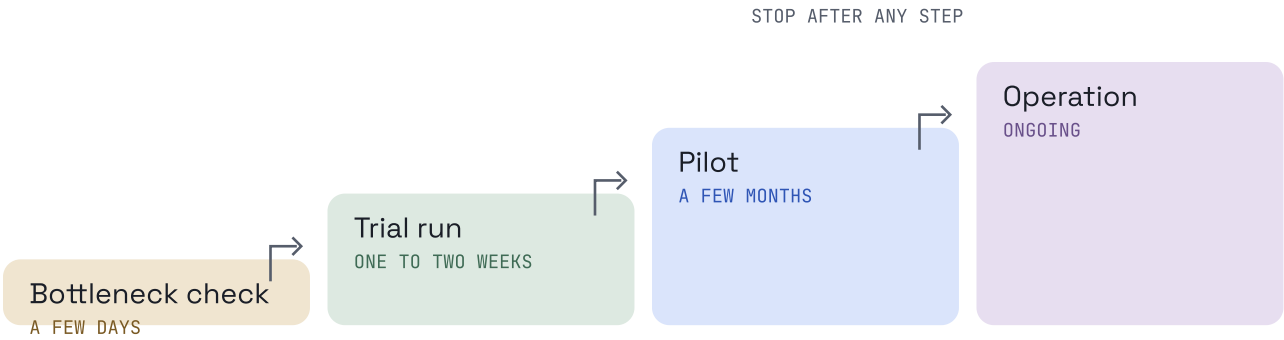
The General Data Protection Regulation (GDPR) gives people the right not to be subject to decisions based solely on automated processing [52]. NIS2 sorts entities into particularly important and important ones, each with its own duties [53]. ISO/IEC 42001 specifies a management system for AI [54], and the Federal Office for Information Security (BSI) describes the risks of generative models for industry and public bodies [55]. When AI is introduced, the works council can call in an expert [56]. Before a pilot processes real documents, its setup is reviewed against these texts with your data protection officer.

We are not lawyers and hold no certificate. Whether a setup complies depends on the use case. The platform is designed to produce the records an auditor can check.

Getting Started

A pilot in four steps, from a bottleneck check to operation.

The first pilot would cover one tender package, from complete plan sets to a signed draft, next to your normal process. Steps three and four need the runtime, which is a concept today. We agree on prices, durations and stop thresholds with you before the first round.



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

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

Released results, signed by named people in your organisation.

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

Glossary

The terms of construction tenders and of the platform used in this paper.

Agent	A model that decides its own next step and which tool to use.
App	One workflow on the platform, with its own screen, agents, rules and connectors.
ATV	Allgemeine Technische Vertragsbedingungen, the technical rules of VOB/C, such as DIN 18299.
Bill of quantities	Leistungsverzeichnis (LV): the structured description of the work in a tender.
Evaluation	A repeatable test of an agent on known cases, with grades and a pass rule.
GAEB	Gemeinsamer Ausschuss Elektronik im Bauwesen, publisher of the exchange format GAEB DA XML.
Gate	A point in the run where a named person signs. The run waits until the signature exists.
Harness	The app where people build, run, check and sign agents.
HOAI	The fee structure for architects and engineers in Germany, with nine work phases.
Kappa	A measure of agreement between two raters that corrects for agreement by chance.
MCP	Model Context Protocol, an open standard to connect models to tools and data.
Ordnungszahl	The number of an element in a bill of quantities, such as 01.02.03.0010.
Realm	The closed space of one organisation on the platform: its data, secrets, models and agents.
Shelves	The registry of the OS from which agents and capabilities install like packages, with a manifest.
VOB	Vergabe- und Vertragsordnung für Bauleistungen, the rules for tendering and contracting construction work.
X83	The GAEB phase for the request for offer, sent from client to bidder.

Sources

56 sources, numbered in the order of first citation. One agent opened each source and a second checked it against the claim, on 24 November 2026. The status stands at each entry.

[1] Destatis: Mehr als jedes dritte Unternehmen nutzt 2026 künstliche Intelligenz. 2026-11-24. https://www.destatis.de/DE/Presse/Pressemitteilungen/2026/11/PD26_408_52911.html *checked in part*

[2] Bitkom: Erstmals nutzt Mehrheit Unternehmen KI. 2026-09-14. https://www.bitkom.org/Presse/Presseinformation/Erstmals-nutzt-die-Mehrheit-der-Unternehmen-KI *checked*

[3] BMWE, prepared by Destatis: Vergabestatistik, Bericht für das Gesamtjahr 2024. 19 Aug 2026. https://www.bundeswirtschaftsministerium.de/Redaktion/DE/Publikationen/Wirtschaft/vergabestatistik-gesamtjahr-2024.pdf *checked*

[4] lxgesetze.de (unofficial consolidated text): VOB/A 2019, § 7. page states 'Zuletzt geändert am 25.9.2023'. https://lxgesetze.de/vob-a-2019/7 *checked, secondary source*

[5] iogentic founding team: Evaluation report, run of 8 November 2026 (internal). 8 November 2026, read 24 November 2026. internal report, available on request *own source*

[6] Yao, Shinn, Razavi, Narasimhan: tau-bench: A Benchmark for Tool-Agent-User Interaction in Real-World Domains (arXiv:2406.12045). 2024-06-17. https://arxiv.org/abs/2406.12045 *checked*

[7] Rabanser, Kapoor, Kirgis, Liu, Utpala, Narayanan: Towards a Science of AI Agent Reliability (arXiv:2602.16666). 2026-02-18. https://arxiv.org/abs/2602.16666 *checked*

[8] Anthropic: Building effective agents. 2024-12-19. https://www.anthropic.com/engineering/building-effective-agents *checked in part*

[9] OpenAI: A practical guide to building agents. 2026-01-15. https://openai.com/business/guides-and-resources/a-practical-guide-to-building-ai-agents/ *checked*

[10] European Commission AI Act Service Desk: AI Act Article 11. 2026-07-27. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-11 *checked*

[11] European Commission AI Act Service Desk: AI Act Article 12. 2026-07-27. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-12 *checked*

[12] European Commission AI Act Service Desk: AI Act Article 14. 2026-07-27. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-14 *checked*

[13] KOFA, IW Köln: KOFA Studie 04/2026, Wertvolle Erfahrung sichern: Ältere Beschäftigte im Unternehmen halten. 2026. https://www.kofa.de/media/Publikationen/Studien/KOFA-Studie_04-2026_AEltere_Beschaeftigte.pdf *checked*

[14] onvista (news agency report on ifo): Ifo: Gut jedes zweite Unternehmen nutzt KI. 2026-06-05. https://www.onvista.de/news/2026/06-05-ifo-gut-jedes-zweite-unternehmen-nutzt-ki-in-der-breite-angekommen-0-20-26518961 *checked, secondary source*

[15] IAB (Institut für Arbeitsmarkt- und Berufsforschung), distributed via idw: IAB press release on generative AI use (IAB-Betriebspanel). 2026-05-05. https://idw-online.de/en/news870242 *checked*

[16] IW Köln, commissioned by Hauptverband der Deutschen Bauindustrie: Gutachten: Bedeutung der Bauwirtschaft. 2024-06-04. https://iwkoeln.de/fileadmin/user_upload/Studien/Gutachten/PDF/2024/Gutachten_2024-Bedeutung-Bauwirtschaft.pdf *checked in part*

[17] Baulinks, reporting a survey by BIngK, VBI, BAK and AHO: Ingenieurbüros unter Druck: Längere Projektlaufzeiten und hoher Bürokratieaufwand. 2026. https://www.baulinks.de/webplugin/2026/1602.php4 *checked, secondary source*

[18] European Commission AI Act Service Desk: AI Act Article 113 (consolidated as at 27 July 2026). 2026-07-27. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-113 *checked*

[19] European Commission: Data Act explained. accessed 2026-11-24. https://digital-strategy.ec.europa.eu/en/factpages/data-act-explained *checked*

[20] Bundesanzeiger Verlag (recht.bund.de): Bundesgesetzblatt Teil I 2025 Nr. 301. 2025-12-05. https://www.recht.bund.de/bgbl/1/2025/301/VO.html *checked in part*

[21] BMJV (gesetze-im-internet.de): BetrVG §87. current text. https://www.gesetze-im-internet.de/betrvg/__87.html *checked*

[22] BMJV (gesetze-im-internet.de): BetrVG §90. current text. https://www.gesetze-im-internet.de/betrvg/__90.html *checked*

[23] DIW Berlin: Bauwirtschaft kehrt zurück auf Wachstumskurs, Wochenbericht 5/2026. 2026. https://www.diw.de/de/diw_01.c.996899.de/publikationen/wochenberichte/2026_05_1/bauwirtschaft_kehrt_zurueck_auf_wachstumskurs.html *checked in part*

[24] Federal Ministry of Justice (gesetze-im-internet.de): § 34 HOAI. HOAI 2021 text, retrieved 24 Nov 2026. https://www.gesetze-im-internet.de/hoai_2013/__34.html *checked*

[25] Statistisches Bundesamt (Destatis): Betriebe, tätige Personen und Umsatz im Baugewerbe. table with reference year 2024, retrieved 24 Nov 2026. https://www.destatis.de/DE/Themen/Branchen-Unternehmen/Bauen/Tabellen/betriebe.html *checked*

[26] Rödl & Partner: Neue EU-Schwellenwerte ab 1.1.2026. 3 Nov 2025. https://www.roedl.com/insights/neue-eu-schwellenwerte-ab-ersten-ersten-zweitausendsechsundzwanzig/ *checked in part, secondary source*

Sources, continued

[27] Federal Ministry of Justice (gesetze-im-internet.de): § 106 GWB Schwellenwerte. in force, retrieved 24 Nov 2026. https://www.gesetze-im-internet.de/gwb/__106.html *checked*

[28] BMJV (gesetze-im-internet.de): GWB §121. current text. https://www.gesetze-im-internet.de/gwb/__121.html *checked*

[29] Institut der deutschen Wirtschaft (IW): Kommune sucht Bauingenieur, IW-Kurzbericht Nr. 41. 17 Jun 2026. https://www.iwkoeln.de/studien/thomas-puls-kommune-sucht-bauingenieur.html *checked in part, secondary source*

[30] Senator für Wirtschaft, Arbeit und Häfen, Bremen: Auswertungsergebnisse der Kostenkontrolle, Vorlage 19/455-L, point 3. 2017-11-09. https://www.wirtschaft.bremen.de/sixcms/media.php/13/19_455_L-Vorlage%20Kostenkontrolle%20Bauprojekte_B.pdf *checked in part*

[31] baunetzwissen.de: Leistungsverzeichnis. undated, read 2026-11-24. https://www.baunetzwissen.de/drucken/leistungsverzeichnis-8494320 *checked in part, secondary source*

[32] tab.de (publisher not verified): Flyer GAEB DA XML. undated. https://www.tab.de/download/554109/GAEB_DA_XML_Flyer.pdf *checked, secondary source*

[33] Flughafen München (operator guideline): Richtlinie zur LV-Erstellung, section 8. 2010-03. https://Munich-Airport.De/_b/0000000000000007428832bb5d8b282c/richtlinie-zur-lv-erstellung.pdf *checked, secondary source*

[34] bauprofessor.de: Neue Generation STLB-Bau. 2020-01-06. https://www.bauprofessor.de/news/neue-generation-stlb-bau/ *checked, secondary source*

[35] DIN Media: DIN 18299 VOB Teil C. 2023-09. https://www.dinmedia.de/de/norm/din-18299/369847346 *checked*

[36] DIN Media: DIN 276 Kosten im Bauwesen. page read 2026-11-24. https://www.dinmedia.de/de/norm/din-276/293154016 *checked*

[37] DIN Media: DIN 18290-2. 2023-11. https://www.dinmedia.de/de/norm/din-18290-2/369464914 *checked*

[38] GAEB: GAEB Datenaustausch. page read 2026-11-24. https://www.gaeb.de/de/produkte/gaeb-datenaustausch/ *checked*

[39] DIN (event BIM nutzt GAEB): BIM nutzt GAEB, Vortrag Datenaustausch. 2017-05-12. https://www.din.de/resource/blob/236706/d25969398ae4db7247ed8a5118ffb2ee/top04-bim-nutzt-gaeb-2017-05-12-vortrag-gaeb-datenaustausch-rheinberg-data.pdf *checked, secondary source*

[40] Niedersächsisches Landesamt für Bau und Liegenschaften: AVA-Pflichtenheft Version 2024-01, section 2.1. 2024-01. https://nlbl.niedersachsen.de/download/208208/AVA-Pflichtenheft_des_SBN_Version_2024-01.pdf *checked in part*

[41] JV Roig: How Much Do LLMs Hallucinate in Document Q&A Scenarios? (arXiv:2603.08274). 2026-03-09. https://arxiv.org/abs/2603.08274 *checked*

[42] Zheng et al.: Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena (arXiv:2306.05685). 2023-06-09. https://arxiv.org/abs/2306.05685 *checked*

[43] Norman, Rivera, Hughes: Reliability without Validity (arXiv:2606.19544). 2026-06-17. https://arxiv.org/abs/2606.19544 *checked*

[44] Lewis et al.: Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks (arXiv:2005.11401). 2020-05-22. https://arxiv.org/abs/2005.11401 *checked*

[45] Anthropic: Introducing Contextual Retrieval. 2024-09-19. https://www.anthropic.com/news/contextual-retrieval *checked*

[46] Google Cloud: A methodical approach to agent evaluation. 2025-11-18. https://cloud.google.com/blog/topics/developers-practitioners/a-methodical-approach-to-agent-evaluation *checked*

[47] Anthropic: Introducing the Model Context Protocol. 2024-11-25. https://www.anthropic.com/news/model-context-protocol *checked*

[48] Model Context Protocol project: Model Context Protocol Specification. 2026-07-28. https://modelcontextprotocol.io/specification/latest *checked*

[49] iogentic: Worked example of the tender package. 24 November 2026. https://iogentic.ai/os *own source*

[50] European Commission AI Act Service Desk: AI Act Article 9. 2026-07-27. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-9 *checked*

[51] EEAS (mirror of Council release): Artificial Intelligence: Council gives final green light to simplify and streamline rules. 2026-06-30. https://eeas.europa.eu/delegations/chile/artificial-intelligence-council-gives-final-green-light-simplify-and-streamline-rules_en *checked*

[52] gdpr-info.eu (unofficial text): GDPR Article 22. n/a. https://gdpr-info.eu/art-22-gdpr/ *checked, secondary source*

Sources, continued

[53] BMJV (gesetze-im-internet.de): BSIG §28. 2025. https://www.gesetze-im-internet.de/bsig_2025/__28.html *checked*

[54] IEC Webstore: ISO/IEC 42001:2023. 2023-12-18. https://webstore.iec.ch/publication/90574 *checked*

[55] BSI: Aktualisierte Fassung veröffentlicht: Chancen und Risiken generativer KI-Modelle. 2025-01-21. https://www.bsi.bund.de/DE/Service-Navi/Presse/Alle-Meldungen-News/Meldungen/Chancen-Risiken-generative-KI_250121.html *checked*

[56] BMJV (gesetze-im-internet.de): BetrVG §80. current text. https://www.gesetze-im-internet.de/betrvg/__80.html *checked*



iogentic

Talk to Us

We offer a thirty-minute conversation without 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

Whitepaper, November 2026. The platform is not released. All photos are generated images. The worked example and the evaluation run use synthetic projects. Sources were read on 24 November 2026.