LAND KÄRNTEN

AI & AUDIT PROCEDURES

AI-Assisted Audit Procedures

From KärntenGPT to a controlled hybrid audit model

Local data sovereignty remains the foundation. External model power is only added after anonymisation.

Mag. Christian Inzko · CIO / CDO
Office of the Carinthian Provincial Government · 5 October 2026

IN THREE STAGES

1 **KärntenGPT**
Sensitive original data stays in Carinthia.

2 **Anonymisation**
Only approved, sufficiently anonymised audit data leaves the protected zone.

3 **External models**
Use complex reasoning where local models are not enough.

Human-in-the-loop
Responsibility stays with the auditor.

Why Carinthia built its own AI platform

Using AI without giving up control over sensitive administrative data.

01

Security & privacy

HR, funding, health or case data should not leave our infrastructure unchecked.

02

Data sovereignty

Models, data storage, logging and access remain under the province's control.

03

Institutional knowledge

Answers should be based on our regulations, documents and business processes – not just on general internet knowledge.

04

Customisability

Use cases and agents can be adapted to Carinthian processes and reused.

OUR STRATEGIC DECISION

Process sensitive original data locally first. Use cloud solutions only where they add clear business value – and only under defined safeguards.

LOCAL FIRST

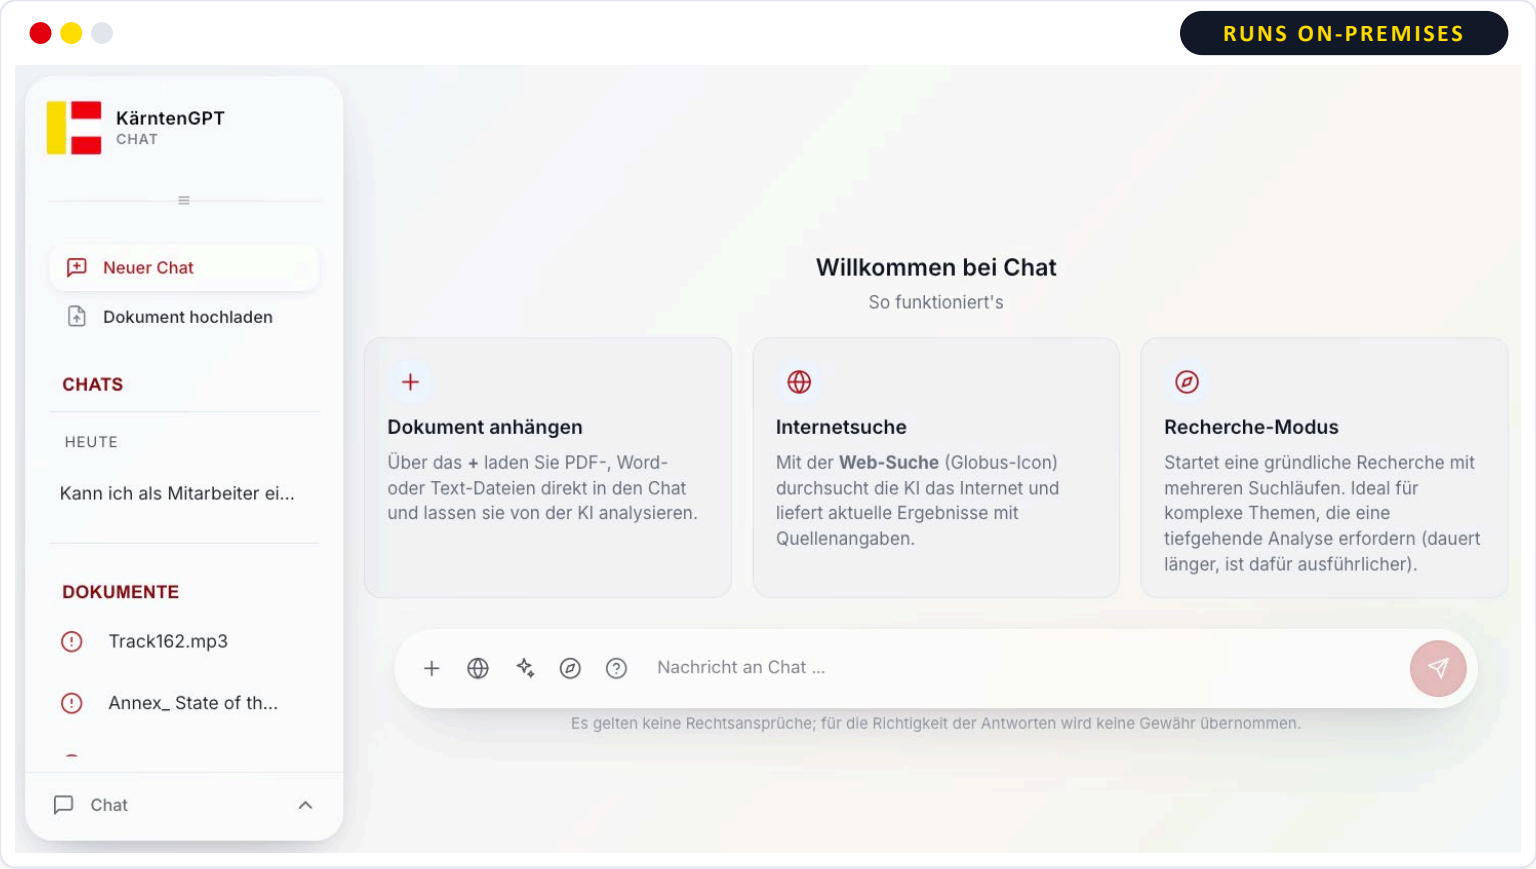
DATA-SOVEREIGN

USE-CASE-DRIVEN

HUMANS DECIDE

Our internal assistant – used like ChatGPT, run on-premises

Staff get familiar access to generative AI without sensitive content being sent to the public internet.



01

A single point of entry

Chat, documents and expert knowledge in one unified interface.

02

Data stays in-house

Sensitive content is processed within our protected zone.

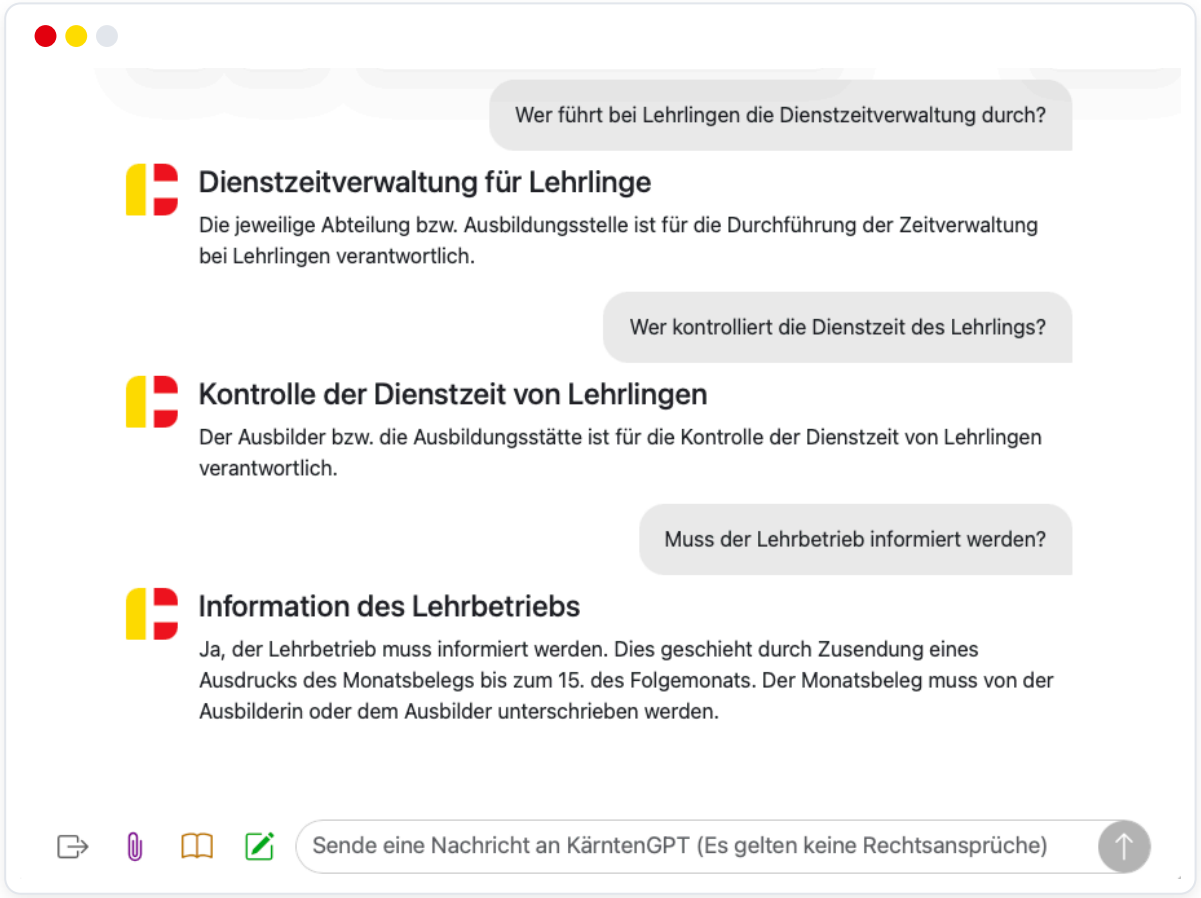
03

A platform, not a single tool

The same foundation supports assistants, document processing and agents.

Internal knowledge makes general AI fit for administration

Regulations, guidelines and expert documents are integrated in a controlled way – without sending them to the public internet.



1

Select the source

Only approved internal documents and regulations are added to the knowledge space.

2

Provide targeted context

KärntenGPT retrieves text passages relevant to the question and provides them to the model as context.

3

Make the answer traceable

The answer is based on the internal knowledge base – not on uncontrolled research.

EXAMPLE

HR questions on working-time rules

The model answers specific questions based on stored regulations and stays within the internal platform.

Example: FAQ editor for staff

Staff can contribute domain-specific knowledge to the local AI themselves via questions and answers.

HOW AN ENTRY IS CREATED

Enter question & answer

Both fields are mandatory.

Classify

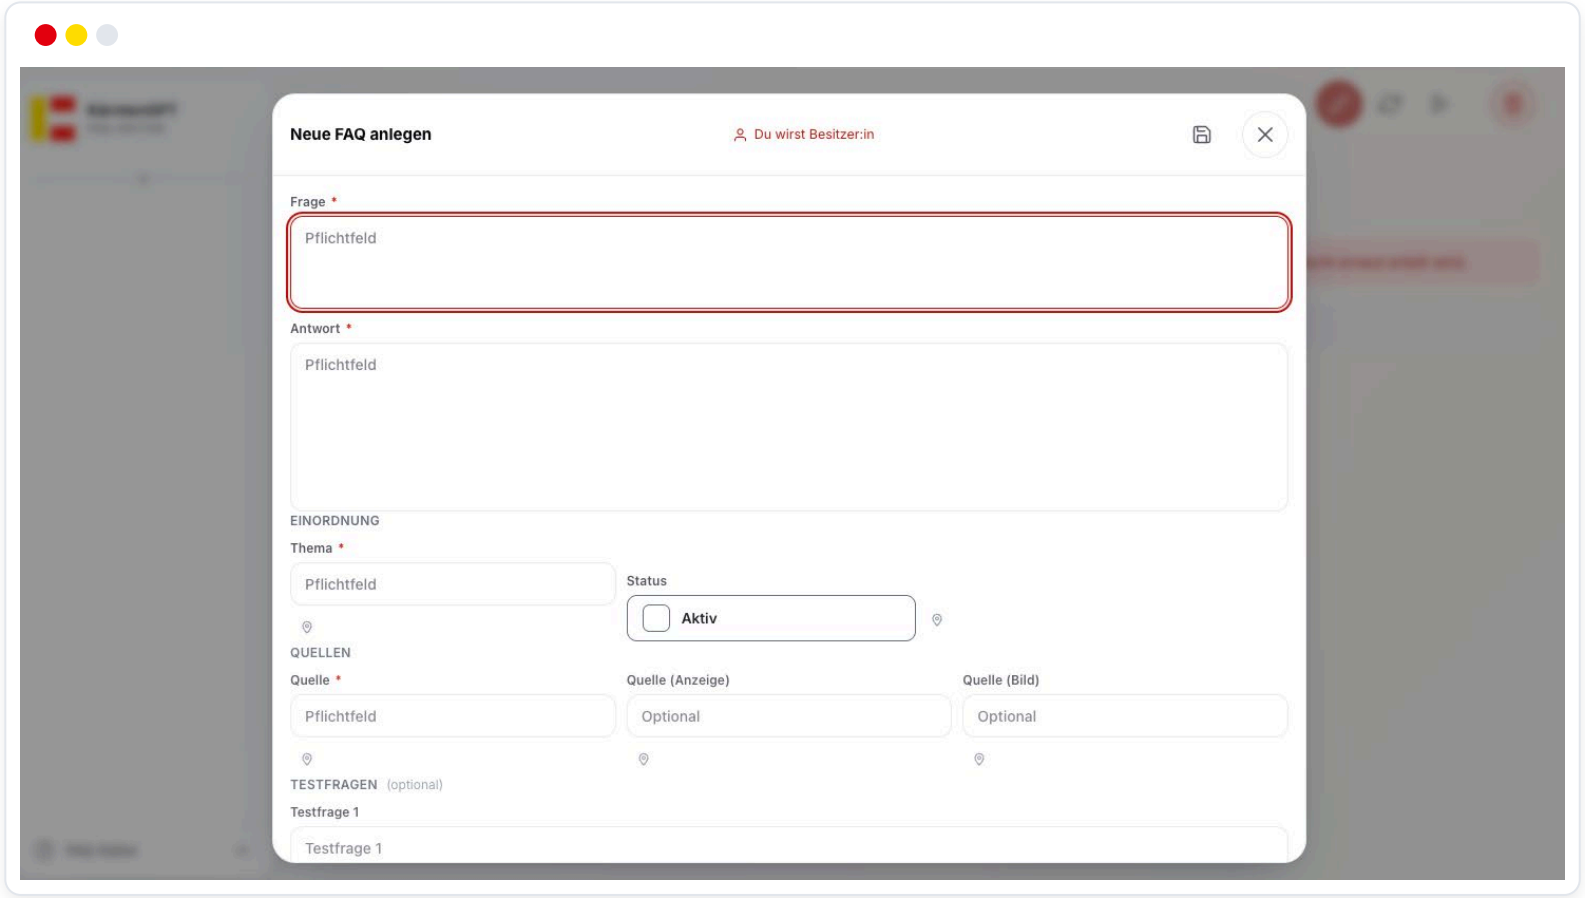
Set topic and status.

Cite the source

Mandatory – display text and image optional.

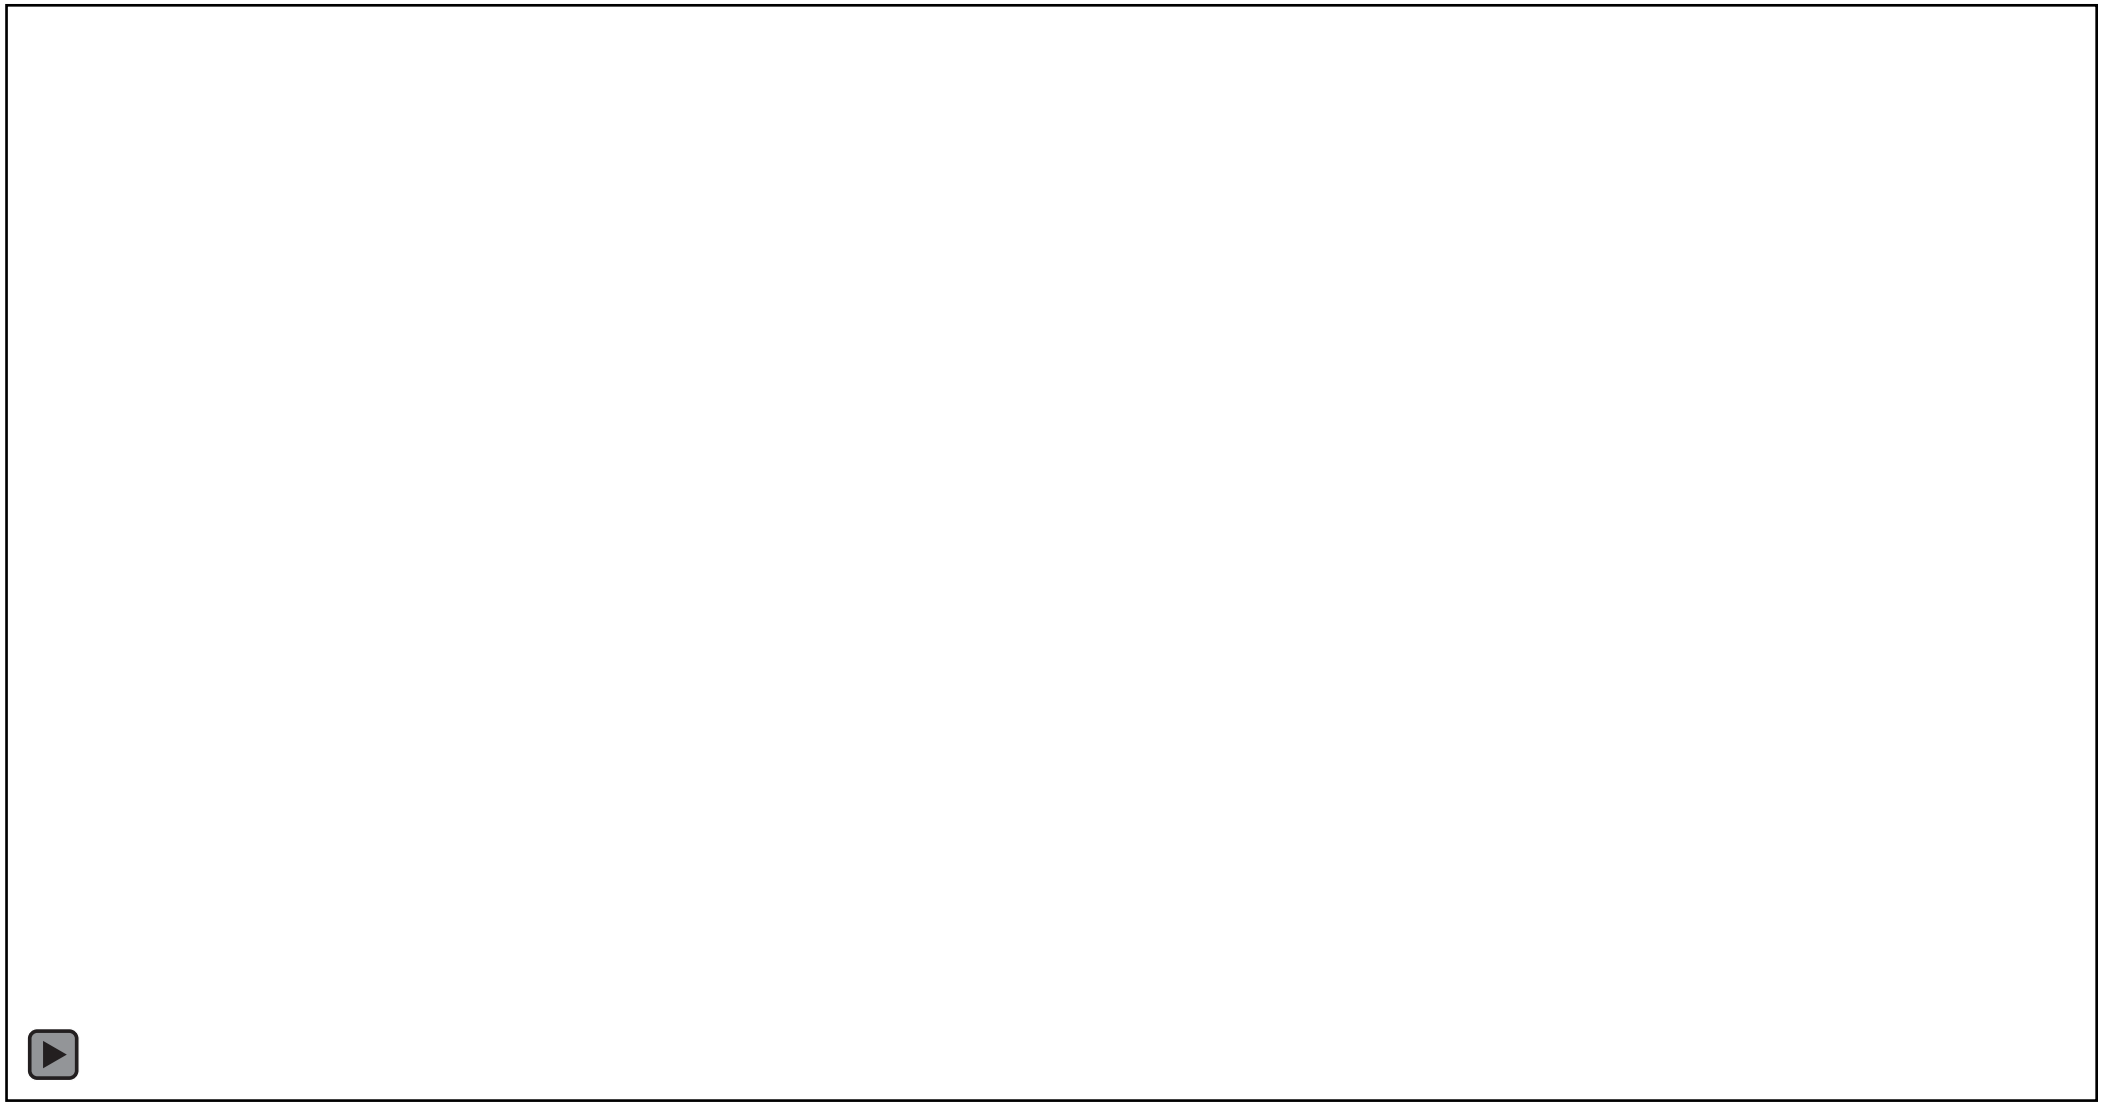
Test questions

Can be added optionally.



KärntenGPT – questions on organisation-specific knowledge

Works like ChatGPT and similar tools – but answers with organisation-specific knowledge



What we learned from local models

KärntenGPT is the right foundation – but “local” is not an end in itself, nor a guarantee of maximum model performance.

Strengths of local models

- **Data protection**
  Original data can be processed internally.
- **Domain knowledge**
  RAG and internal sources are controllable.
- **Integration**
  Direct connection to internal processes and systems.
- **Cost control**
  No usage-based cloud costs for every standard case.

vs.

Limits with demanding tasks

- **Complex reasoning**
  Multi-step audit chains are harder for small local models.
- **Long context**
  Many contracts, documents and time series at once strain context and quality.
- **Multimodality**
  Tables, scans and heterogeneous documents require strong models.
- **Stability of structured outputs**
  Robust, consistent results are particularly important for audit processes.

CONSEQUENCE **KärntenGPT remains the secure control point – external models complement it selectively.**

What KärntenGPT cost

Started externally in 2023 – then know-how and further development were deliberately brought in-house.

TOTAL INVESTMENT

approx. €140,000

Development + hardware – excluding ongoing internal staff costs

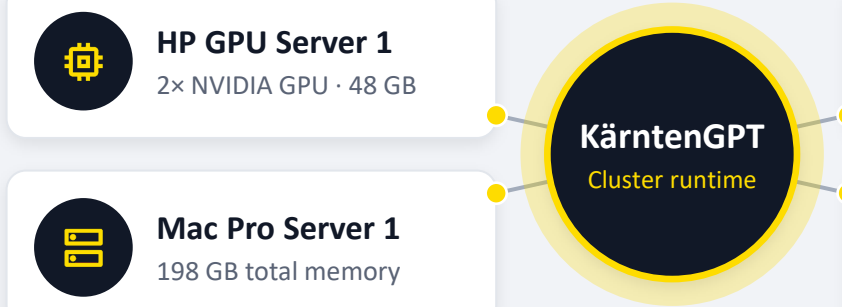


€50,000
Initial external development (2023)

€90,000
On-premises AI cluster (hardware)

ON-PREMISES AI CLUSTER 4 servers · 4 NVIDIA GPUs · 2 Mac Pro

Hardware **approx. €90,000**



All four systems work as one shared compute cluster. Workloads can be distributed between GPU servers and Mac Pro nodes.

2023 · Started externally
Initial build of KärntenGPT by an external company.

Then · Know-how brought in-house
Taken over by newly hired in-house staff.

Today · In-house operation
Operation, integration and further development are handled in-house.

STRATEGY

Accelerate externally at first – master know-how, operations and infrastructure in-house for good.

Why Court of Audit use cases are particularly demanding

The AI must not just read – it must connect information across multiple sources.

01

Many sources

Contract, purchase order, invoice, proof of service, approvals, correspondence and time data must be considered together.

02

Many rules

Daily rates, surcharges, caps, deadlines, service periods and formal mandatory details apply simultaneously.

03

Many relationships

A deviation often only becomes visible when several documents, periods or people are compared.

04

High traceability

Every audit flag must be traceable to specific data or document passages – no black-box decisions.

WHAT THE AI SHOULD DO

Structure documents

Detect deviations

Compare time series

GOAL

Auditors spend more time evaluating – and less time searching, sorting and comparing.

Where KärntenGPT alone can reach its limits

Not every audit step needs an external model. But some tasks benefit significantly from greater model performance.

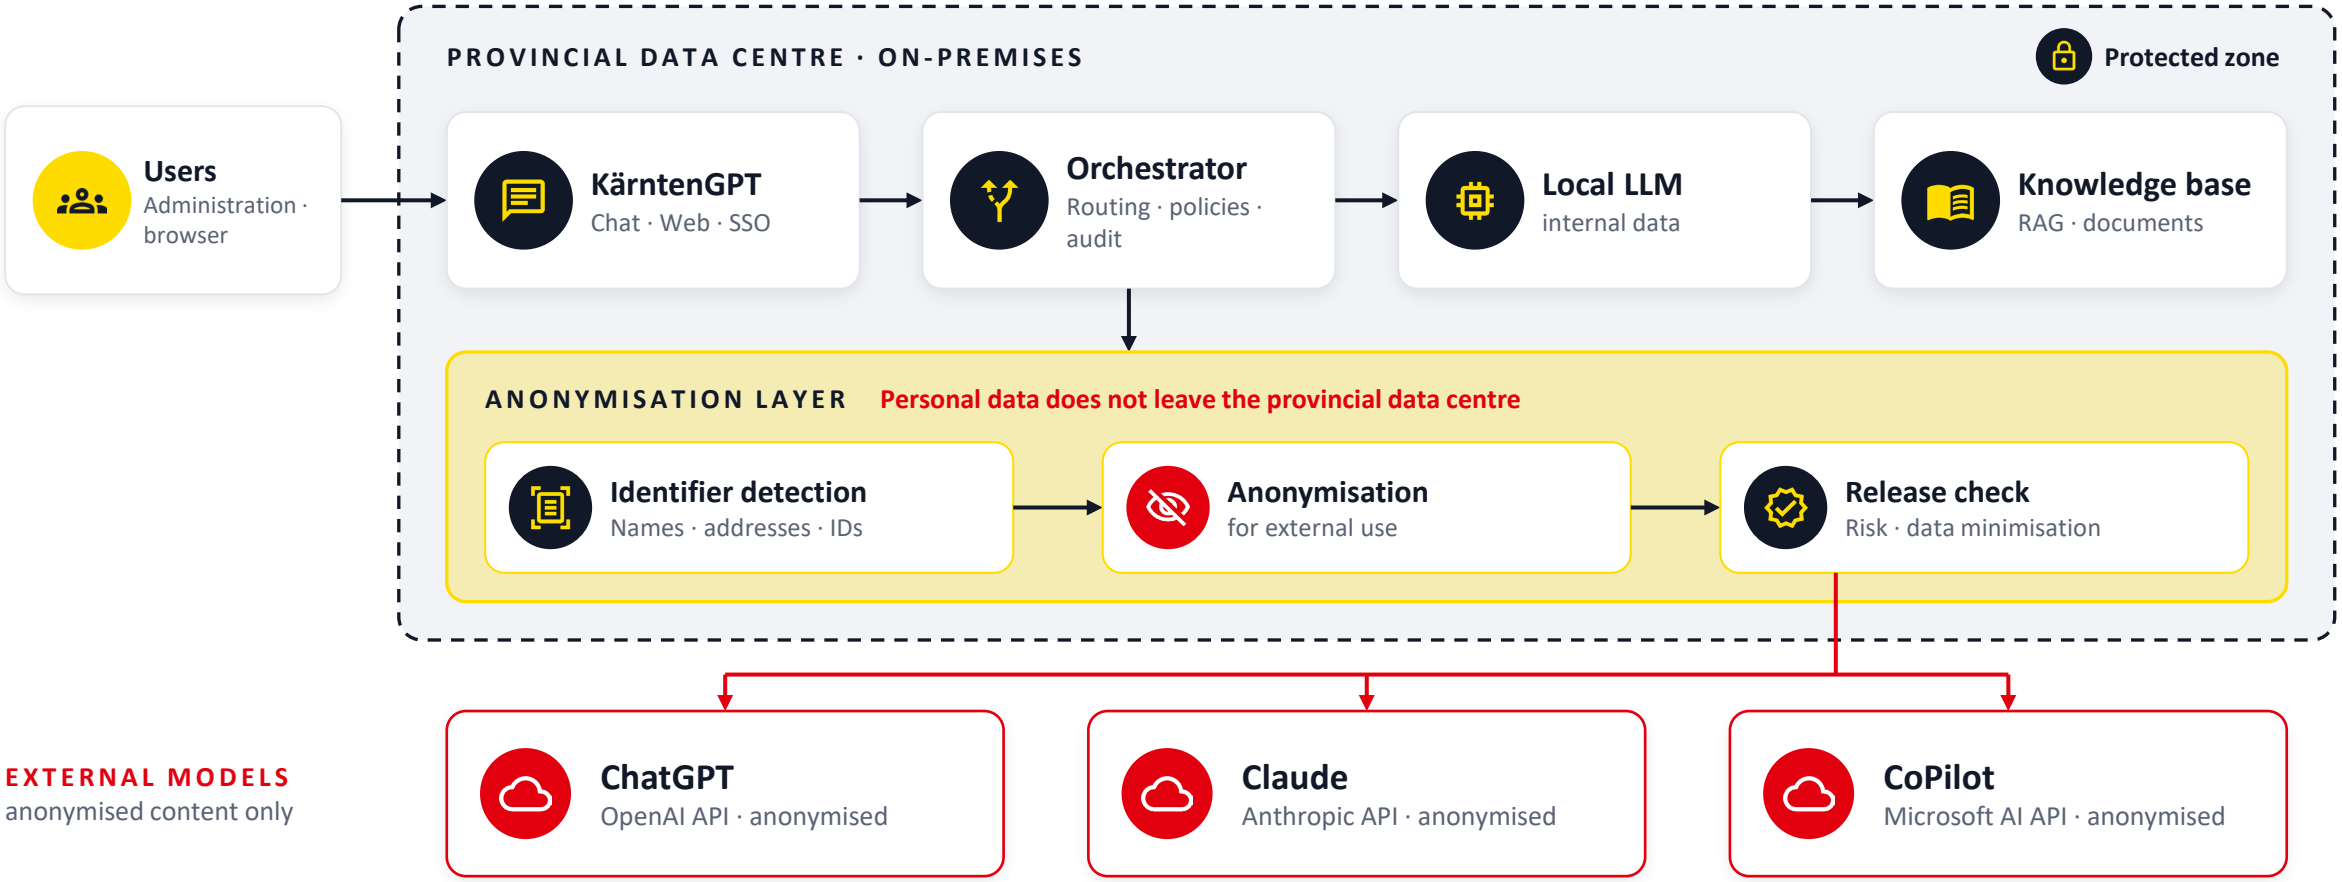
TASK	KÄRNTENGPT LOCAL	LARGE EXTERNAL MODEL	ROUTING
Document classification & extraction	●●● **very good**	●●○ good	**prefer local**
Internal knowledge queries / RAG	●●● **very good**	●●○ good	**prefer local**
Several long contracts + invoices at once	●○○ limited	●●● **very good**	**consider external**
Complex expert justification / reasoning	●●○ good	●●● **very good**	**consider external**
Unstructured tables / scans / mixed formats	●●○ good	●●● **very good**	**consider external**
Outlier and pattern analysis with extensive context	●●○ good	●●● **very good**	**consider external**



SECURITY **KärntenGPT orchestrates model selection. The external model never gets direct access to electronic records (ELAK), SAP or original documents.**

KärntenGPT 3.0 – the province's secure AI hub

Local data sovereignty remains the foundation. External frontier models are only connected in a controlled way and exclusively with anonymised content.



Anonymisation ≠ pseudonymisation

Both reduce the link to individuals – but with different consequences for analysis and re-identifiability.

Anonymisation

The link to individuals is removed so that re-identification is no longer possible with reasonable means.

EXAMPLE

“[PERSON] has 18 overtime hours.”

The person can no longer be linked across multiple documents.

≠

Pseudonymisation

Direct identifiers are replaced by a pseudonym. Re-linking remains possible via separately protected additional information.

EXAMPLE

“PERSON_017 has 18 overtime hours.”

Cross-references across documents are preserved – without transmitting the name.

RULE OF THUMB

- Identity analytically irrelevant → **anonymise**
- Cross-references required → **pseudonymise**

One record – three levels of data protection

Same business content, but a different link to individuals.

ORIGINAL · INTERNAL

Name	**Anna Berger**
Staff no.	**4711**
Department	**IT Operations**
Invoice	**RE-2026-0831**
Service	**31 Aug · 1 PD**

→ **maximum analytical freedom**

ANONYMISED

Name	**[REMOVED]**
Staff no.	**[REMOVED]**
Department	**IT Operations**
Invoice	**RE-2026-0831**
Service	**31 Aug · 1 PD**

→ **lowest re-identification risk**

PSEUDONYMISED

Name	**PERSON_017**
Staff no.	**ID_017**
Department	**IT Operations**
Invoice	**RE-2026-0831**
Service	**31 Aug · 1 PD**

→ **links remain auditable**

AI detects unusual overtime patterns in a municipality

The AI flags anomalies – the auditor clarifies context, approval and factual justification.

DATA BASIS

- 36-month time series
- Target and actual hours
- Overtime & time off in lieu
- approved / not approved
- Department & employment level

pseudonymised before external AI

LLM
Pattern analysis

AUTOMATED OUTLIER DETECTION

HIGH	**MA-017** continuous increase		**+186%**
HIGH	**MA-042** persistently high level		**avg. 29 h**
MEDIUM	**MA-031** sudden peak		**47 h**
POSITIVE	**MA-056** steady reduction		**–62%**

Audit flag – not an automatic accusation or decision. Trends are illustrative.

Why we need local AND external

Routing by protection needs, data status and analytical complexity.

ROUTING

KärntenGPT orchestrates model selection

	USE CASE	MODEL	DATA STATUS	OPERATION
A	**Sensitive standard cases and processing of original data**	KärntenGPT / internal LLM	Original data & sensitive documents	On-premises
B	**Time series, cross-references and cross-document patterns**	External LLM + pseudonymisation	Pseudonyms instead of direct identifiers	Cloud / API
C	**Demanding expert analysis where identity is irrelevant (contracts)**	External LLM + anonymisation	No meaningful link to individuals	Cloud / API

PRINCIPLE **Controlled routing: local where possible – external model power where it adds business value.**

What an audit-proof AI review process requires

Technology alone is not enough. Approvals, logging, source references and clear responsibilities are decisive.

1 DATA → 2 MODEL → 3 RESULT

Data approval

Which data may be processed locally? Which may be used externally – once anonymised?

Model routing

For each audit step it is defined which model may be used and why.

Source references

Every relevant AI flag must trace back to data points or document passages.

Anonymisation rules

Automated detection plus defined thresholds and spot checks.

Prompt & version management

Audit rules, prompts and model versions must be documented and traceable.

Human-in-the-loop

No autonomous payment, HR or legal decisions by the model.

TARGET STATE **AI as a traceable audit assistant – not an autonomous decision-maker.**

IN PRACTICE

A practical example

From a pseudonymised overtime dataset to a traceable audit finding

LIVE DEMO / DEEP DIVE

PROCESS

1 **Data**
pseudonymised

2 **Flag**
Patterns & outliers

3 **Decision**
by the auditor

Source file – example payroll data

ORIGINAL DATA · INTERNAL Direct identifiers such as name, date of birth, social security number and address are still included.

PERSNR	NACHN	VORNA	TITEL	GESCH	GBDAT	SVNR	FAMST	ANZKD	STAAN	STRAS
Personalnummer	Nachname	Vorname	Titel	Geschlecht	Geburtsdatum	Sozialvers.-Nr.	Familienstand	Anz. Kinder	Staatsangehörigkeit	Straße
00010001	Schuh	Luca	Dr.	1	23.07.1975	9053230775	02	2	AT	Am Kanal 159/20
00010006	Barbieri	Andrea	Dr.	2	26.06.1983	5036260683	01	0	AT	Mühlweg 103
00010010	Prasser	Katharina	Mag.	2	25.06.1987	7132250687	02	1	AT	Buchenweg 209/21
00010015	Willegger	Markus	Prof.	1	02.05.1965	6564020565	02	2	AT	St. Veiter Straße 27/38
00010019	Grill	Gudrun	Mag.	2	22.04.1992	7421220492	01	0	AT	Am Bach 173
00010024	Steinberger	Oezge	BEd	2	10.07.1999	2772100799	04	1	IT	Burggasse 45/9
00010029	Lang	Sofia		2	17.05.1978	9077170578	01	2	AT	Kastanienallee 214/19
00010033	Hollersbacher	Luana	BSc	2	30.04.1973	8866300473	01	0	HR	Wörthersee-Straße 35
00010038	Tomic	Hans		1	03.07.1979	5960030779	01	1	AT	Hangweg 2
00010039	Aktas	Sibel		2	24.06.1966	3102240666	02	1	AT	Tannenweg 30
00010043	Preiml	Biljana		2	05.09.1997	6663050997	01	0	AT	Talstraße 125
00010048	Grbic	Ana	Dipl.-Ing.(FH)	2	21.12.1979	2953211279	01	2	HR	Hauptplatz 87
00010050	Kroell	Valentina		2	20.04.1981	7789200481	01	2	SI	Hangweg 143
00010052	Kljajić	Erwin	MBA	1	09.03.1976	1644090376	01	0	HR	Rosentaler Straße 104

PRACTICAL EXAMPLE · LIVE DEMO

Anonymisation / pseudonymisation process

Step 1 · Data is pseudonymised before external analysis.

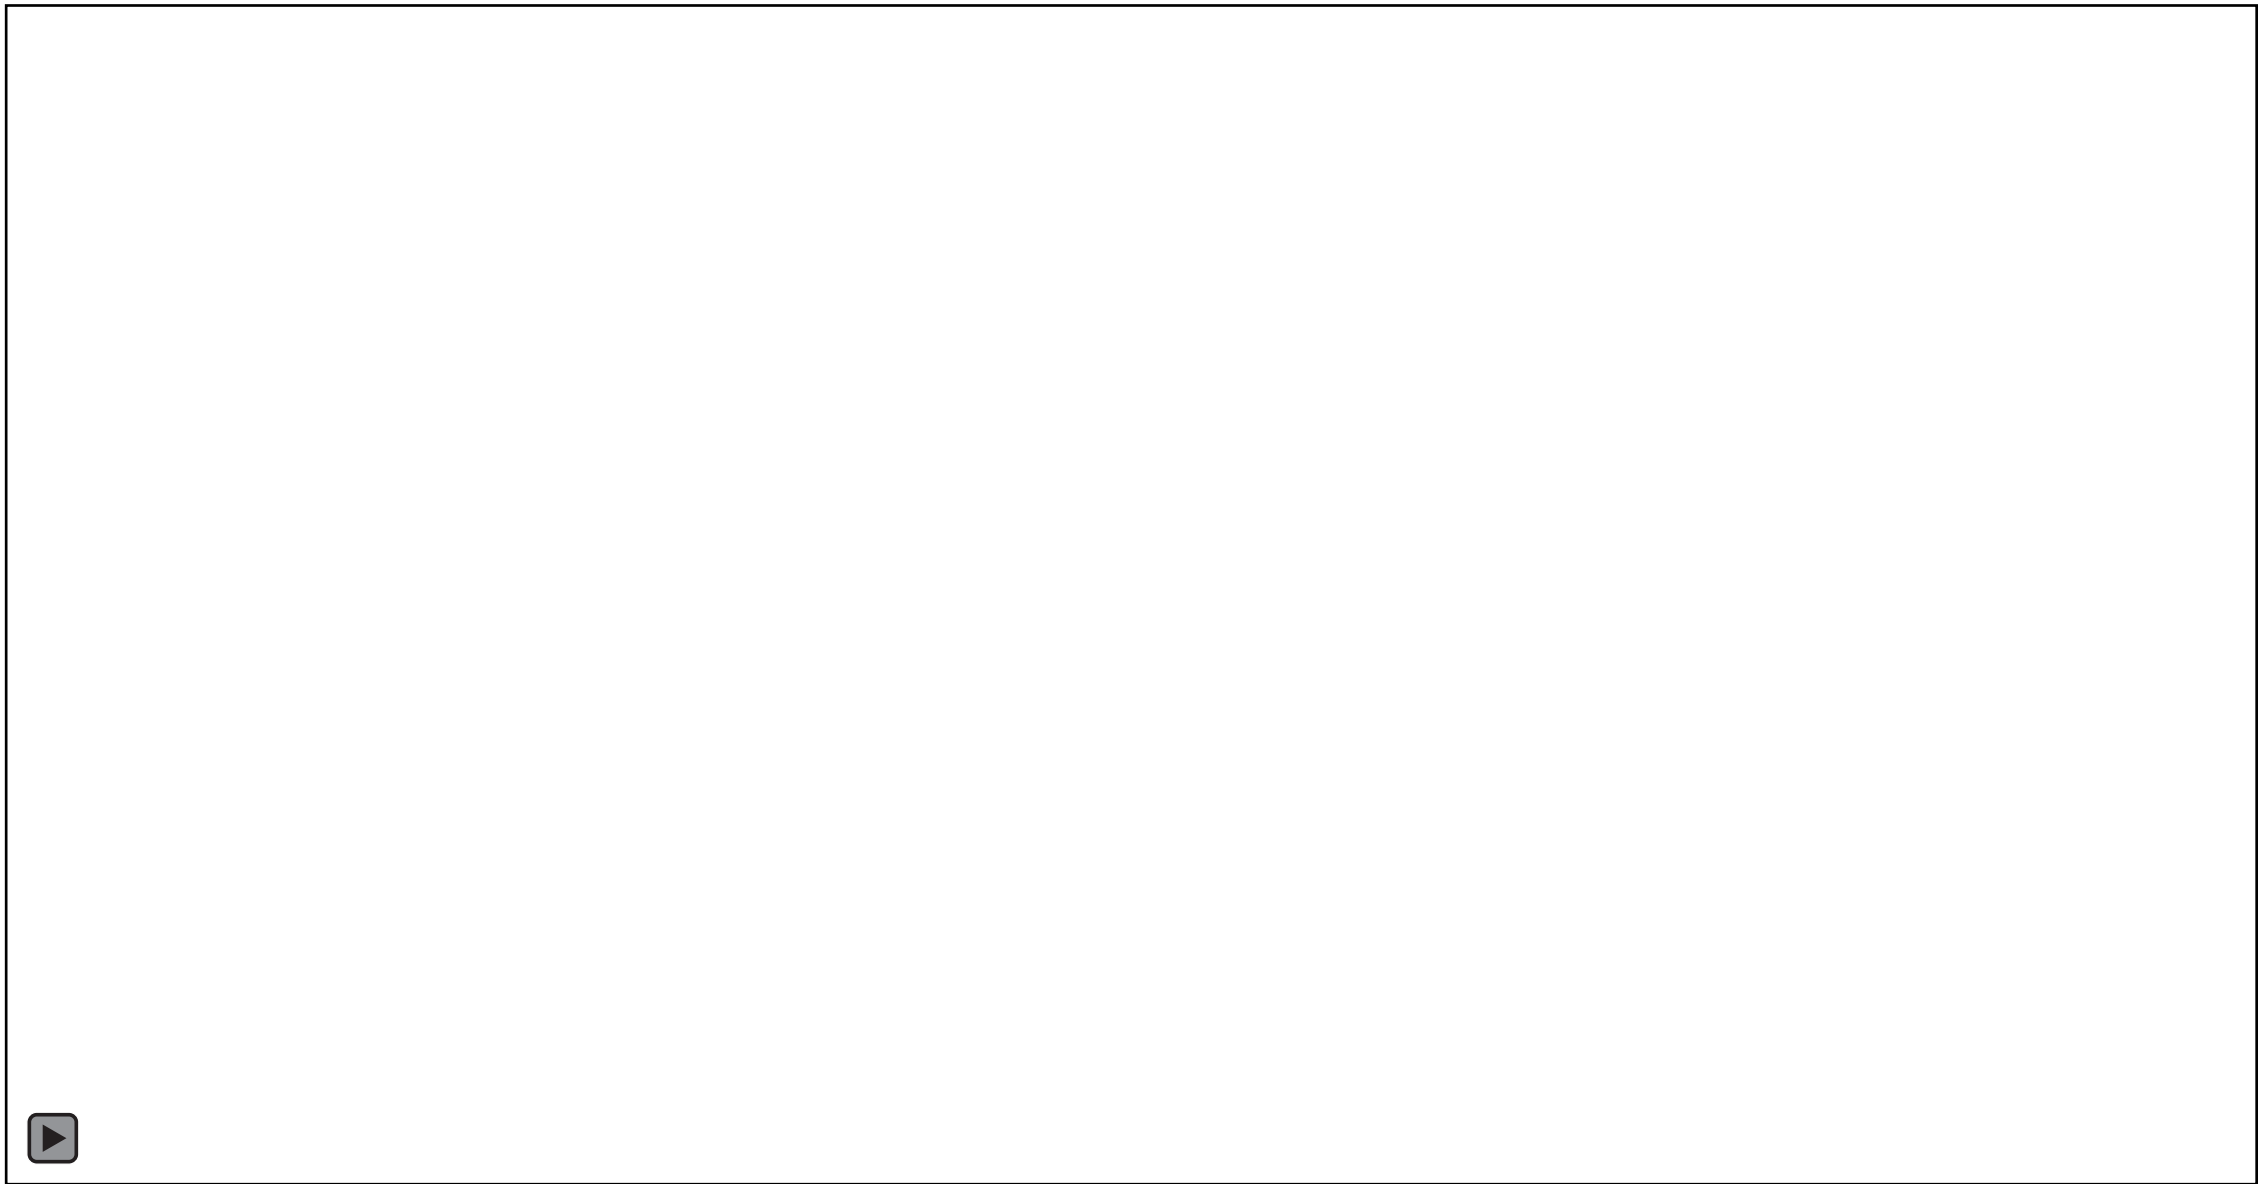
PROCESS

1 **Data**
pseudonymised

2 **Flag**
Patterns & outliers

3 **Decision**
by the auditor

Anonymisation / pseudonymisation process



PRACTICAL EXAMPLE · LIVE DEMO

The result …

Step 1 · The pseudonymised dataset

PROCESS

1 **Data**
pseudonymised

2 **Flag**
Patterns & outliers

3 **Decision**
by the auditor

Anonymisation / pseudonymisation process



PRACTICAL EXAMPLE · LIVE DEMO

AI analysis

Step 2 · Detect patterns & outliers

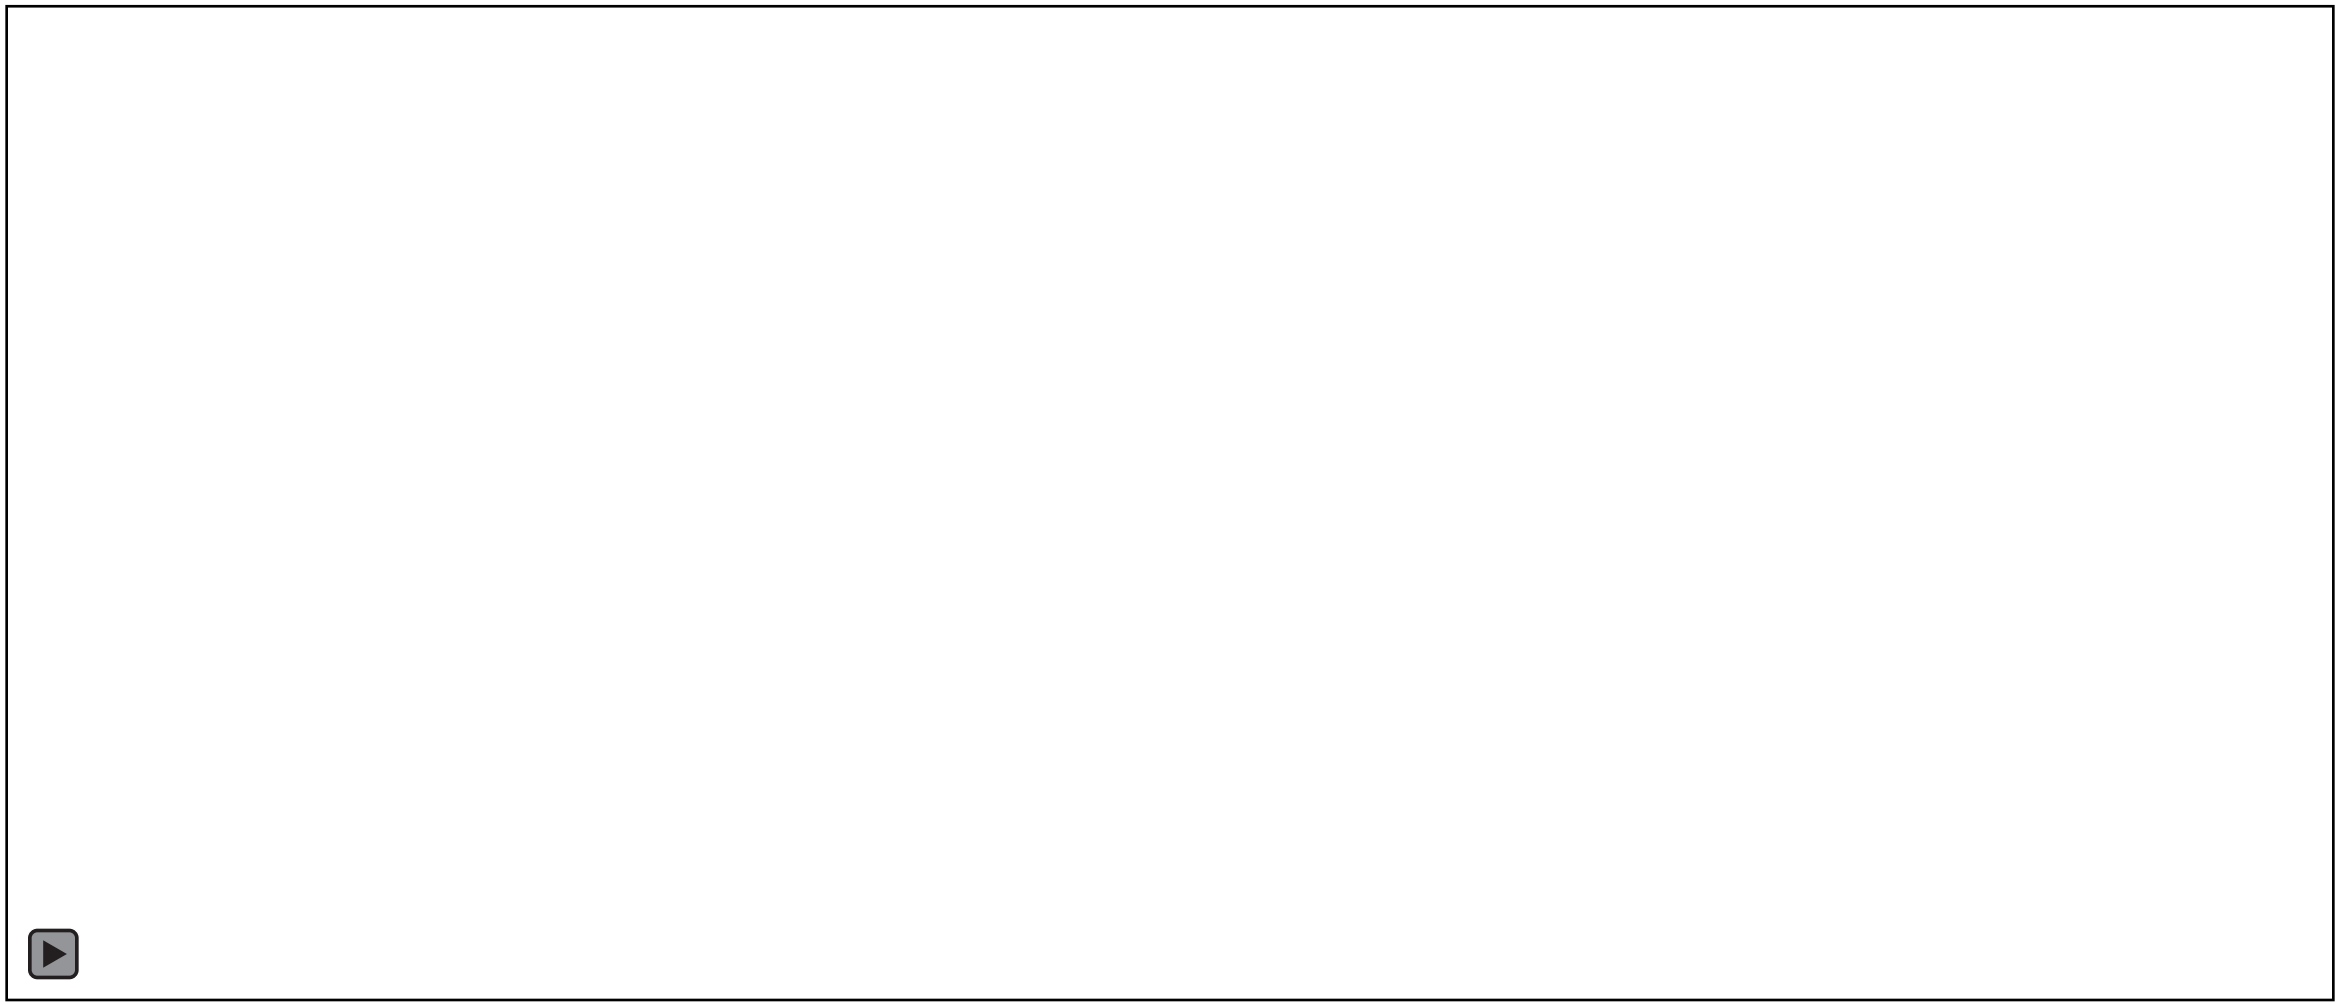
PROCESS

1 **Data**
pseudonymised

2 **Flag**
Patterns & outliers

3 **Decision**
by the auditor

AI analysis



AI analysis

What AI does not solve automatically

Good AI auditing requires data quality, source references and clear escalation rules.

RISK		CONTROL
Hallucinations	A model can produce plausible but unsupported statements.	**Accept only source-referenced, structured findings.**
Incomplete data	Missing amendments or approvals can lead to wrong conclusions.	**Explicitly flag missing documents as uncertainty.**
Ambiguous contracts	Legally ambiguous clauses cannot be resolved by technical means alone.	**Escalation to subject-matter or legal review.**
Pseudonym leakage	Too many contextual attributes can enable indirect re-identification.	**Data minimisation, thresholds and approval rules before external use.**

What we gain from the hybrid model

The architecture combines the strengths of local data processing with the power of large models – without abandoning the data protection principle.

Less manual work

Supporting documents and data are automatically consolidated, structured and pre-checked.

Greater completeness

Rules and comparison patterns are applied systematically to all relevant items.

Risks visible earlier

Price, quantity, time or approval deviations are specifically prioritised.

Better traceability

Audit flags come back structured, justified and with source references.

THE KEY ADDED VALUE

Not “more AI”, but better routing: simple and sensitive tasks locally, complex analysis externally – but only on anonymised data.

CONCLUSION

Controlled routing instead of “local only” or “cloud only”.

1

KärntenGPT remains the foundation

Original data, internal knowledge, orchestration and logging remain in our infrastructure.

2

External models are added selectively

We use them where complex reasoning and large context windows add measurable value.

3

External means anonymised

No personal original or pseudonymised data leaves the protected zone.

4

Humans decide

AI prioritises, explains and supports – audit responsibility stays with the auditor.

Thank you for your attention

Mag. Christian Inzko · CIO / CDO
Office of the Carinthian Provincial Government