


ASK KARGENTRA · GOVERNED AI

Ask Kargentra

Ask Kargentra, the logistics AI assistant that knows your operation

Ask Kargentra is the logistics AI assistant built into every signed-in Kargentra screen. It reads the record in front of you, answers with the facts it used and where they came from, and proposes the next permitted step. It never acts on its own: a person with the right permission approves through the same governed action, and the audit trail records both.



ASK KARGENTRA · DISPATCH READINESS · CITED ANSWERS, PROPOSALS ONLY

CONTROL TOWER · OPERATIONS · EVERY QUEUE THAT NEEDS A PERSON, MOST URGENT FIRST

WHAT CHANGES FOR THE OPERATION

Answers you can check

Every fact in an answer carries its source and freshness, and unknowns are named rather than guessed.

Proposals, not surprises

Ask Kargentra suggests the next permitted step with a reason and an owner, and a person approves it.

Context without AI

The Context Layer works with AI switched off: facts, blockers and next steps are deterministic.

Governed per tenant

Each AI use case is switched on or off per tenant, and every invocation is recorded as evidence.

WHAT ASK KARGENTRA DOES

Ask on every screen

One launcher on every signed-in operations, warehouse, supplier, carrier, driver and admin screen, aware of the screen and its record.

- Ask about this screen, explain this screen, summarise this record
- Screen descriptors kept complete by an automated coverage check
- Context disclosure: permitted record facts, or screen guidance only

Always-on Context Layer

A deterministic layer states what each screen is for, which record it binds and what is blocking it, with or without AI.

- Provenance and freshness on every fact
- Freshness unknown stated plainly when it is
- Recent changes read from canonical events

Cited answers

Answers are grounded only in records the asking user is already permitted to read, with each fact cited.

- Permission-checked grounding on every request
- Named unknowns instead of invented values
- A failing source becomes a stated refusal, not a wrong answer

Proposals a person approves

Suggested next steps come with a reason, an owner and a link to the permitted workflow, never as a hidden button.

- Advisory until a person approves
- Unsupported or unauthorised actions shown as refused
- Approval runs through the governed action and the audit trail

Module assistants

Purpose-built prompts in each module: orders, planning, dispatch, exceptions, warehouse, supplier, returns and driver help.

- Orders: summarise, why blocked, explain the shortfall, what happens downstream
- Planning: explain these options, checks before acceptance
- Warehouse and supplier: explain, prioritise and draft

Drafts for people to send

Customer updates, conversation replies and return responses are drafted for review, and a person decides whether to send them.

- Every draft marked as not sent
- Internal notes never quoted in summaries
- Price, liability, claims and credits handed to a person

HOW IT FLOWS

How a question becomes an approved action

- 01 Govern**
A tenant administrator decides which use cases are on, within the platform envelope. With AI off, the Context Layer still answers.
- 02 Ask in context**
A user asks from the screen they are on. The Context Layer supplies the screen's purpose and the record it binds.
- 03 Check permission**
Grounding reads only what this user may already read, in this tenant and within their record scope.
- 04 Answer with citations**
The answer states each fact with its source and freshness, and names whatever is unknown.
- 05 Propose**
A useful next step is proposed with a reason, an owner and a link to the permitted workflow. Its state: proposed, not applied.
- 06 Approve or decline**
A person with the right permission opens the governed action and decides. Nothing changes until they do.
- 07 Record**
The invocation and the human decision are recorded, so the audit trail shows what was suggested and who acted.
- 08 Improve**
Users flag answers as outdated, wrong route or missing step, and reviewed feedback becomes a canonical answer.

ASK KARGENTRA IN ASK KARGENTRA

"What can I do here?"

TRIP-4410 is Blocked on two checks. The carrier's insurance certificate on file expired on 30 September, and the warehouse hand-over is not complete: 4 of 6 handling units for SHP-30781 are staged and 2 are still at Pack. The driver, tractor and trailer assignment is eligible, and there are no open exceptions on the movement.

READINESS · CARRIER INSURANCE · EXPIRED 30 SEP

WAREHOUSE · SHP-30781 · 4 OF 6 UNITS STAGED

ASSIGNMENT · DRIVER, TRACTOR, TRAILER · ELIGIBLE

EXCEPTIONS · NONE OPEN

PROPOSAL · APPROVED BY THE DISPATCHER

Request an updated insurance certificate in the carrier conversation, then re-evaluate readiness once the last two units are staged.

Ask Kargentra reads only what your role can open, cites every fact and proposes. A named person approves; nothing is sent or executed on its own.

EXCEPTIONS AND EDGE CASES HANDLED

- A user asks about a record outside their scope: the answer is refused, and nothing about that record is disclosed.
- One grounding source fails: that source is reported as refused and the rest of the answer stands on its own facts.
- AI is switched off for the tenant: the Context Layer still shows the complete facts and says that it does.
- A user asks Ask Kargentra to release a trip: the action is shown as refused, with a link to the dispatcher's workflow.
- A conversation touches price, liability or a claim: the topic is handed to a person instead of being drafted.
- A fact's freshness cannot be established: the answer says Freshness unknown rather than presenting it as current.

CAPABILITY REFERENCE

Where it works

- Operations, warehouse, supplier, carrier and driver apps
- Tenant and platform administration
- Customer portal with customer-safe grounding
- Carrier portal with carrier confidentiality
- Phone, tablet and desktop

Grounding and citation

- Permission-checked reads only
- Shipment, exception, return, account, invoice, dispute and conversation records
- Provenance and freshness on every fact
- Named unknowns
- Per-source refusal on failure

What it does

- Explains a screen or a record
- Summarises and diagnoses blockers
- Suggests next steps with reason and owner
- Drafts replies and updates for people to send
- Answers what you can do with your access

What it never does

- Write to operational records
- Release, award, accept or cancel
- Send anything outside Kargentra
- Decide a price, claim, credit or dispute
- Read beyond the user's permissions

Governance

- Tenant AI control per use case
- Eleven governed use cases
- Platform-level AI envelope
- Recorded invocation evidence
- Deterministic answers first; model synthesis optional
- Event-triggered briefs under tenant policy

STANDARDS AND CONNECTIONS

DETERMINISTIC ANSWER ENGINE

TENANT-CONFIGURED MODEL PROVIDER

PERMISSION-CHECKED READ MODELS

CANONICAL EVENT TRIGGERS

AUDIT EVIDENCE STORE

QUESTIONS BUYERS ASK

What can Ask Kargentra answer?

Ask Kargentra answers questions about the screen and record in front of you: what it is, why it is blocked, what changed recently, what the options mean and what to do next. It also answers what you can do with your access and how your workflow runs. Every answer is built from records you are already permitted to read, and cites them.

Does Ask Kargentra take actions on its own?

No. Ask Kargentra explains, cites and proposes. It has no write access, and actions such as release, award, acceptance, cancellation or sending a message are shown as refused if asked for. A proposal links to the permitted workflow, where a person with the right permission decides, and both the suggestion and the decision are recorded.



See Ask Kargentra on your own scenario

Bring a real lane, site or exception. More at www.kargentra.com/products/ask-kargentra/

info@kargentra.com

WWW.KARGENTRA.COM · SIGN IN: APP.KARGENTRA.COM