

Statement of services

Statement of services

The subject of the contract is the one-off creation of a language model from the documents supplied by the customer, together with its digital provision through an activation code and a setup program for download. The basis is an open base model of the Qwen3.5 family under the Apache 2.0 licence, adapted by LoRA fine-tuning: the pre-trained weights remain unchanged and only additional matrices are trained. The customer runs the result on its own, self-procured hardware inside its own network.

STATUS 12 August 2026 PUBLISHER IonKon GmbH

1. Subject of the service

After delivery, IonKon processes no customer content any longer. The delivered system has no channel back to IonKon and runs without an internet connection once it is set up. One order produces exactly one model from exactly the material uploaded for that order.

Expressly not part of the service:

- Hardware. IonKon supplies no device and brokers none.
- An ongoing service, a maintenance contract or a cloud connection.
- Operation, support or monitoring of the model inside the customer network.
- A later fine-tuning run with new material; that requires a new order.
- Guarantees in the legal sense. Descriptions on the website and in this document are not guarantees.

2. Key figures

All prices are net and exclude statutory VAT. The price is charged once; no recurring licence or usage fees arise, and the number of users inside the customer network is not limited.

The text volume is the governing cap, not the file size: what is counted are characters of extracted text, converted at 1,000 characters per page of business text. The ZIP limit is enforced technically at upload. The dataset generated for fine-tuning is always the same size regardless of how much you upload; the volume of your material is therefore the boundary, not the reason for choosing a tier.

The graphics memory figure covers the model weights only. Every concurrent conversation adds a cache that grows with the length of the history. How to work out the requirement for your own number of users is set out in [What hardware an AI model needs in house](#).

Item	Value
Conclusion of contract	on receipt of the order confirmation by e-mail
Data processing agreement	concluded no later than before the data upload
Advance payment	none
Invoice	at the start of processing, full model price, net, payable within 14 days
Processing time	tier M usually a few days, tier L about a week; dates are approximate unless expressly committed as binding
Provision	activation code by e-mail to the address held in the account
Download window	30 days from provision
Deletion after provision	once you confirm receipt, and no later than 30 days after provision
Deletion without provision	no later than 180 days from the start of the order

Item	Value
Deletion record	issued automatically once deletion has taken place
Notification of a personal data breach	without undue delay, no later than 48 hours after becoming aware
Warranty period	12 months from provision
Audit rights	once a year and whenever there is a specific reason; on-site inspection with 14 days' notice
Change of a sub-processor	announced in advance in text form, objection on serious grounds within 14 days
Deliverables, model	weights of the fine-tuned model, ownership certificate with SHA-256 checksum, Apache 2.0 licence text and NOTICE files
Deliverables, software	setup program for Windows, Linux and macOS, AnythingLLM interface for browser access inside the company network
Deliverables, documents	verification tool for the marking of outputs, short guide to checking it, this service description

3. Sequence and timeline

The sequence is the same for both tiers. The timeline starts with the order and ends with the deletion record; the only step that materially affects the duration is your upload, because processing starts only after it.

You need internet access only for sign-in, activation and the single download. Operation afterwards runs entirely inside the company network. While work is under way you can see the state of processing in your account, broken down into checking your material, creating the training data, fine-tuning and provision.

Step	What you do	What IonKon does
Order	choose a tier, describe the company and the use case, accept the terms, confirm business status	receive the order
Conclusion of contract	receive the order confirmation	send the order confirmation by e-mail, open the order portal
Data processing	conclude the data processing agreement	provide the agreement with the technical and organisational measures as its annex
Upload	upload the documents as a single ZIP file in the portal	issue a short-lived presigned upload form, validate the archive on the server
Start of processing	receive the invoice, net, 14 days	issue the invoice, create a dedicated GPU instance for this order
Preparation	follow progress in your account	extract text, remove personal data, generate training examples
Fine-tuning	no involvement required	train the model, check quality, store the package encrypted, destroy instance and volumes
Provision	take the activation code from the e-mail	send the activation code, keep the package available for download for 30 days
Setup	download the setup program, sign in with account and code, install the model on your own hardware	provide sign-in, activation and download
Confirmation of receipt	confirm provision via the link sent to you	send two reminders if confirmation does not arrive

Step	What you do	What IonKon does
Deletion	receive the deletion record	delete uploaded documents, training data and IonKon copies of the model, issue the record
Operation	run the model in your own network, offline	no further processing

4. What you provide

The following points sit with the customer. They are duties to cooperate, not recommendations: without them the service cannot be performed or the result cannot be put into operation.

- **Hardware:** a server or workstation with a supported NVIDIA graphics card and a current driver. Cards of the Blackwell generation use the faster FP4 mode, all other supported cards run in BF16 mode automatically, which quadruples the memory requirement given in section 2.
- **Software on the target system:** Docker with Compose v2 and the NVIDIA Container Toolkit. The setup program checks both and offers a one-click installation if Docker is missing.
- **Network:** the company network for your team's browser access, plus internet access for sign-in, activation and the single download. On macOS additionally SSH access to a separate Windows or Linux server with a suitable graphics card, because Mac machines have no NVIDIA card.
- **People:** one person with administrator rights on the target system, an IonKon account with e-mail and password, and access to the e-mail address on file, to which the activation code is sent. For the selection of documents, one person responsible for the subject matter.
- **Documents:** a ZIP file with formats that carry a text layer, that is Word, PowerPoint, Excel with text content, text-based PDF, Markdown, HTML and TXT. Scanned PDFs without a text layer, photographs, CAD drawings, video and audio yield no text and only consume space in the archive. Folder structure and file names are irrelevant.
- **Rights in the content:** the assurance that you are entitled to use the uploaded content for fine-tuning and that it contains no unlawful material. Special categories of personal data

under Article 9 GDPR are excluded by contract.

- Responsibility after delivery: storage and backup of the downloaded copy. It is unaffected by any deletion at IonKon and rests with you alone.

5. What you decide

Four decisions rest with the customer, two of them before the order. They determine the price, the date of the invoice, the content of the model and the permitted scope of use.

- Tier: M or L. What decides is not the volume of your material alone but the task. M covers one clearly bounded field with consistent language, L covers several fields, inconsistent or contradictory sources, several languages and tasks spanning several steps. That capability comes with model size and cannot be substituted by more training data.
- Timing: the invoice goes out when processing starts, and processing starts after your upload. You therefore control both the invoice date and the start of the 30-day download window.
- Volume and selection of documents: the yardstick is what the model should know, not what happens to be at hand. Old versions, duplicates and bare columns of figures without explanatory text consume your allowance to no benefit.
- Intended purpose: an internal assistance system based on your own company documents. Use for practices prohibited under the AI Regulation, or as a high-risk system under Annex III, requires prior agreement with IonKon.
- Interface: the supplied interface, in which the notice about the AI system is already set up and permanently visible, or an interface of your own in which you implement that notice yourself.
- Optional details at the order stage: a name under which the model presents itself and a tone that shapes the style of all answers. Without them it presents itself neutrally as your company's assistant.

6. Rights in the model, contract, payment

On payment in full, all transferable rights in the fine-tuned model, including the weights, pass irrevocably to the customer. IonKon retains no rights of use, exploitation or recall. Until

payment in full, a simple, revocable right of use applies; default in payment does not affect the provision of a model that has already been completed.

The customer may use, modify, reproduce and commercially exploit the model without limitation in time, territory or content, including resale. The only terms to observe are the Apache 2.0 conditions of the base model, essentially the licence notice; licence text and NOTICE files accompany the provision. Proof is the ownership certificate carrying the SHA-256 checksum of the weights.

The offering is addressed exclusively to entrepreneurs within the meaning of section 14 of the German Civil Code, to legal entities under public law and to special funds under public law, not to consumers. German law applies, excluding the UN Convention on Contracts for the International Sale of Goods, and the place of jurisdiction is the registered office of IonKon GmbH. Amendments and additions require text form. The model and its software may be subject to European and US export control law; the customer is responsible for compliance on any transfer or re-export.

No advance payment is charged. At the start of processing IonKon issues an invoice for the full model price, net plus statutory VAT, payable within 14 days. In the event of default, statutory default interest under section 288(2) of the German Civil Code applies. The customer may set off only claims that are undisputed or have been finally determined by a court.

As an intangible asset acquired for consideration, the model can in principle be capitalised and depreciated; the amount to recognise is the acquisition cost, that is the price paid. The hardware remains a separate asset. Please settle the specific accounting treatment with your tax adviser; the full three-year calculation is set out in [What an in-house AI model really costs](#).

7. Data protection and deletion periods

For the uploaded documents the customer is the controller and IonKon is the processor under Article 28 GDPR; for the master data of the customer account IonKon is its own controller. The upload travels through a short-lived presigned form directly into the encrypted object storage of STRATO GmbH in Germany and never touches the web servers. Content processing runs on a GPU instance created for that order alone at DataCrunch Oy, cloud brand Verda, in Finland, or at Akenes SA, brand Exoscale, in Frankfurt, Zagreb or Zurich, and deleted afterwards; which of them depends on available capacity. BuchhaltungsButler GmbH in Germany receives master data only for invoicing and payment reconciliation, and no training documents. All four

process within the European Union or Switzerland; no external AI provider gains access to customer documents at any point.

Two deletion periods apply. Where provision takes place, uploaded documents, the training data derived from them and the model copies remaining at IonKon are deleted once you confirm receipt, and no later than 30 days after provision. Where provision does not take place, for instance because an order fails technically, deletion follows no later than 180 days from the start of the order. Once deletion has taken place, a deletion record is issued automatically. A local copy you have already downloaded is unaffected. Order and invoice data are retained within the commercial and tax retention periods under section 257 of the Commercial Code and section 147 of the Fiscal Code.

Before the training examples are generated, personal data is removed automatically from the extracted text. That reduces the risk of such data remaining in the finished model but does not eliminate it: IonKon's own impact assessment rates precisely this risk as medium to high. The duty to carry out a data protection impact assessment falls on the controller, which means the customer; IonKon supports this under Article 28(3)(f) GDPR and provides its own assessment as a basis. The individual stages, the processors and the residual risk are set out in [Where your documents go when a model is built](#).

8. Roles under the AI Regulation

IonKon is the provider of the delivered system within the meaning of Regulation (EU) 2024/1689 and in that role fulfils the marking obligation under Article 50(2): every output issued through the interface carries machine-readable, signed metadata, and a supplied tool verifies the signature independently of IonKon. LoRA fine-tuning does not reach the compute threshold above which an adaptation would itself become a base model with its own cascade of provider obligations. The allocation of roles depends on the circumstances of the individual case.

The customer is the deployer. It falls to the customer to give the notice to interacting persons under Article 50(1), which is already configured in the supplied interface, to keep to the intended purpose, to promote AI literacy among staff under Article 4, and to seek prior agreement if a prohibited or a high-risk use under Annex III is planned. Which obligations do not apply to an internal assistance system, and which deadlines are running, is set out in [The AI Regulation: what applies to ordinary companies and what does not](#).

9. Limits of the service

The following points are neither owed nor warranted. They appear here so that expectations of the result are settled before the order rather than negotiated after delivery.

- Correctness of individual outputs. AI models generate outputs mechanically on a statistical basis; no particular correctness, completeness or fitness is owed. Outputs must be checked before business-critical use.
- Freedom from false statements and from verbatim reproduction of training content. The state of the art allows neither to be promised; what helps organisationally is set out in [What a model learns from your records, and what it does not](#).
- A particular quality of result. It depends substantially on the volume and quality of the documents supplied. If the data is unusable, for instance because files are damaged or carry no text layer, you are given the opportunity to supply suitable data within a reasonable period.
- Up-to-date knowledge. The state of knowledge is the state of the uploaded material. For individual facts that change frequently, the document retrieval in the supplied interface is intended, which you maintain yourself without further fine-tuning: [Search documents or train a model](#).
- Marking outside the interface. The marking attaches to the system's response, not to the text. If the plain text is copied into an e-mail or another system, the marking does not travel with that copy.
- Operation on undersized hardware. If the weights do not fit into graphics memory, the system falls back on main memory and becomes slower by orders of magnitude. There is no such thing as a device that is only slightly too small.
- Claims arising from delay or from quality. Only where provision fails for reasons for which IonKon is responsible does the claim to remuneration lapse and are amounts already paid refunded.
- Legal advice. This document describes the state of the service and does not replace a legal review of your own case.

10. Status and validity

Status: 12 August 2026. The key figures in section 2 reflect the tiers and the legal texts as at that date. Prices are subject to change and are net plus statutory VAT.

The binding documents are the general terms and conditions, the data processing agreement including its annex on technical and organisational measures, and the privacy policy. In the event of any discrepancy those documents prevail, and in matters of data protection the data processing agreement does. Acceptance of the terms is stored with its timestamp and text version; an acceptance already given remains valid with the version that was accepted.

This document is updated as soon as tiers, legal texts or the technical implementation change. The version published on ionkon.de always governs.

Sources

[IonKon AGB](#)

[IonKon Auftragsverarbeitungsvertrag](#)

[IonKon Technische und organisatorische Maßnahmen](#)

[IonKon Datenschutzerklärung](#)

[IonKon Preise](#)

[IonKon KI-Verordnung](#)

[IonKon Einrichtungsprogramm](#)

[IonKon Modelle und Hardware](#)

[Welche Hardware ein KI-Modell im Haus braucht](#)

[Wohin Ihre Dokumente gehen, wenn ein Modell entsteht](#)

[Was ein Modell aus Ihren Unterlagen lernt, und was nicht](#)

[KI-Verordnung: was für normale Firmen gilt und was nicht](#)

[Was ein eigenes KI-Modell wirklich kostet](#)

[Dokumente durchsuchen oder Modell anlernen](#)

[Verordnung \(EU\) 2016/679 \(DSGVO\), EUR-Lex](#)

[Verordnung \(EU\) 2024/1689 \(KI-Verordnung\), EUR-Lex](#)

[Hu et al., LoRA: Low-Rank Adaptation of Large Language Models, 2021](#)