

WHITEPAPER



ChatGPT®, Claude®, Gemini®: The best AI is in the US. Your data stays 100% anonymous. 100% GDPR-compliant.

Anonix Privacy Proxy
Use AI. Protect data. No compromises.

Version 1.1 – October 2026

Convecto GmbH

Contents



The Invisible Risk	2
The CLOUD Act	3
What Really Happens When Your Software Uses AI APIs	4
The Solution: Anonix Privacy Proxy	5
How the Protection Works	6
Government, Public Administration & Public Sector	7
Hospitals and Clinics	8
Care Facilities	9
Courts and Justice	10
Social Services and Youth Welfare Offices	11
Law Firms	12
Tax Advisors and Auditors	13
Banks and Insurance Companies	14
HR Departments	15
Educational Institutions	16
Pharma and Life Sciences	17
Software Companies and IT Service Providers	18
Why Existing Solutions Fall Short	19
The Architecture at a Glance	20
Investment and Cost-Effectiveness	21
Compliance at the Push of a Button	22
Voices from the Field	23
Your Path to Anonix	24
Legal Notice	25

The Invisible Risk



Imagine this: your line-of-business application sends citizen data to an AI service via API to make official notices easier to understand. Your law firm software has draft contracts containing client names reviewed. Your hospital system asks the AI about treatment options and transmits the patient record in the process. Or an employee copies sensitive data straight into ChatGPT®.

It happens. Every day. In thousands of German companies and public authorities. And in most cases, management has no idea.

54%

OF KNOWLEDGE WORKERS
use unofficial AI tools at work

42%

OF COMPANIES
assume their employees use AI
privately

49%

OF EMPLOYEES
would keep using AI tools even if
they were explicitly banned

Sources: Bitkom Research 2025, Cornerstone OnDemand 2025

Why Bans Don't Work

Companies that simply ban AI lose twice. First, employees use the tools anyway—just without oversight and without any safeguards. Second, these companies give up the productivity edge that AI tools provide. Market research firm Gartner predicts that by 2030, 40 percent of all organizations worldwide will experience security incidents caused by uncontrolled AI use.

The real question is not whether your employees use AI. It's whether you know which data leaves your organization when they do.

The CLOUD Act—the Problem No One Sees

OpenAI® is based in San Francisco. So is Anthropic®. Google®, of course. What many decision-makers don't know: the US CLOUD Act of 2018 gives American authorities the right to demand that US companies hand over stored data. No matter which server the data is on. No matter which country.

When your software or your employees send personal data to OpenAI®, an OpenAI®-compatible provider such as DeepSeek® or Mistral®, to Claude® or Gemini®—whether through an API integration or direct use—US authorities can, in theory, access that data. This is in direct conflict with the GDPR.

US CLOUD Act

US authorities may demand data from US companies worldwide

Applies regardless of where the data is physically stored

No judicial review by EU standards required

Affects OpenAI®, Google®, Microsoft®, Anthropic®, DeepSeek®, Mistral®, xAI®, Groq®, Perplexity®, and all other US and cloud providers

VS

EU GDPR

Personal data requires a legal basis for transfers to third countries

The Schrems II ruling invalidated the Privacy Shield
The EU-US Data Privacy Framework applies only to certified companies

Art. 48 GDPR prohibits disclosure without a mutual legal assistance treaty

The Supposed Solution: An EU Data Center

Many companies believe they have solved the problem by not booking AI models directly from the provider, but running them through AWS® Bedrock or Microsoft® Azure® in a European data center. The data never physically leaves the EU—problem solved?

No. The CLOUD Act is tied not to the location of the server but to where the company is headquartered. AWS® is a subsidiary of Amazon® (USA). Microsoft® Azure® is a product of Microsoft® Corporation (USA). Both are subject to the CLOUD Act—regardless of whether the data is processed in Frankfurt, Dublin, or Stockholm.

Legal reality: An EU data center run by a US provider does not protect against the CLOUD Act. US authorities can demand the data, and the provider is legally required to cooperate. This has been confirmed by both the European Data Protection Board (EDPB) and the CJEU's Schrems II ruling.

With Anonix, the server location no longer matters: the AI provider—whether accessed directly or through an EU data center—sees only anonymized data with placeholders. Even in the event of a CLOUD Act request, there is nothing usable to hand over.

The Consequences Are Real

5.9

BILLION EUR

in GDPR fines across Europe since 2018

1.2

BILLION EUR

largest single fine (Meta®, 2023, for third-country transfers)

62%

OF COMPANIES

in Germany transfer data to non-EU countries

Sources: DLA Piper GDPR Fines Survey 2025, Bitkom Research 2025

The dilemma: Using AI without data protection is illegal. But forgoing AI is economically reckless. The solution has to deliver both.

What Really Happens When Your Software Uses AI APIs



Let's look at a typical process. Your line-of-business application calls an AI API to make an official notice easier to understand:


API request to the LLM provider:

Please rewrite the following notice so it is easier to understand:


Dear Ms. Ingrid Bergmann,
Your application for housing benefit (Ref. WG-2026-04817) has been approved.
Payment account: DE89 3704 0044 0532 0130 00
Address: Goethestr. 14, 60313 Frankfurt am Main

Without Anonix, four categories of personal data leave your infrastructure: name, case reference, IBAN, and full address. They end up on servers in the US.

BEFORE: Without Protection



AFTER: With Anonix Privacy Proxy



What the LLM provider sees:

Please rewrite the following notice so it is easier to understand:

Dear **[PERSON_1]**,
Your application for housing benefit (Ref. **[CASE_NO_1]**) has been approved.
Payment account: **[IBAN_1]**
Address: **[ADDRESS_1]**, **[ZIP_1]**

The AI works just as well with the placeholders as with the real data. In the response, Anonix automatically reinserts the original data—in real time, even for streamed responses. Nothing changes in the workflow for your software or your employees.

The Solution—Anonix Privacy Proxy

Anonix is a transparent proxy that sits between your applications and the AI provider—whether OpenAI®, DeepSeek®, Mistral®, Groq®, Perplexity®, xAI® Grok®, Claude®, Gemini®, or a self-hosted model. Every request is automatically anonymized before it leaves the EU infrastructure. Every response is automatically de-anonymized before it reaches the employee.

Zero integration: Simply swap the API URL and API key in your existing software. No rebuild. No code changes. Anonix works with any software that uses OpenAI®-compatible, Claude®, or Gemini® APIs—whether in-house business applications, SaaS products, or developer tools.

What Makes Anonix Unique

For Decision-Makers

- **GDPR-compliant by design**—personal data never leaves the EU
- **CLOUD Act neutralized**—US providers see only placeholders
- **No loss of productivity**—employees work as usual
- **EU-hosted SaaS**—ready to use immediately, full data sovereignty
- **Multi-tenant**—isolated data per department or customer

For IT

- **Cloud-based**—no server of your own required, ready to go immediately
- **All major LLM APIs**—OpenAI® + all OpenAI®-compatible (DeepSeek®, Mistral®, Grok®, and many more), Claude®, Gemini®
- **Streaming**—real-time de-anonymization, even for SSE responses
- **Admin GUI**—manage providers, configure rules, monitor costs
- **Audit log**—complete logging, CSV export

Supported AI Providers and Models


PROVIDER / PROTOCOL	MODELS (EXAMPLES)	STREAMING
OpenAI® and all compatible APIs	GPT®-4o, GPT®-4, GPT®-3.5 Turbo, o1, o3	Yes (SSE)
OpenAI®-compatible providers	DeepSeek®, Mistral®, xAI® Grok®, Perplexity®, Together AI®, Groq®, OpenRouter®, and many more	Yes (SSE)
Anthropic® Claude®	Claude® 4 Opus, Claude® 3.5 Sonnet, Claude® 3 Haiku	Yes (SSE)
Google® Gemini®	Gemini® 2.0 Flash, Gemini® 1.5 Pro, Gemini® 1.5 Flash	Yes (SSE)
Local / self-hosted models	Ollama®, LM Studio®, vLLM® (any OpenAI®-compatible model)	Yes (SSE)

The key point: Any software that uses an OpenAI®-compatible, Claude®, or Gemini® API works with Anonix. That includes hundreds of providers—from OpenAI® to DeepSeek®, Mistral®, xAI® Grok®, Perplexity®, and Groq® all the way to self-hosted models. Whether in-house business applications, ERP systems, law firm software, developer tools (Cursor®, GitHub® Copilot®), AI frameworks (LangChain®, LlamaIndex®), or direct browser use.

How the Protection Works



Anonix doesn't detect personal data with a single method, but through four successive layers of protection. Each layer catches what the previous one missed.



Layer 1: Pattern Recognition

More than 60 patterns detect structured data such as email addresses, phone numbers, IBANs, tax IDs, monetary amounts, and dates. This works reliably and quickly because these data types follow a fixed format.

Layer 2: Anonix ConvNet NER

For unstructured data such as personal names, company names, and places, a neural network is used. It understands linguistic context and recognizes that "Munich" is a place and "Ms. Weber" is a name. In eight languages: German, English, French, Spanish, Italian, Dutch, Portuguese, and Polish.

Layer 3: Anonix Zero-Shot NER

Some industries have their own sensitive data types: diagnoses, case references, project numbers, cost centers. Anonix detects these without any programming. You describe the category in your own words, and the system learns it instantly.

Layer 4: AI Quality Check

A dedicated small language model checks the already anonymized text once more for overlooked fragments. It acts as a second opinion that judges independently of all other layers.

Before and After: An Example

Original text:

Dr. Thomas Richter (t.richter@kanzlei-richter.de) represents Ms. Ingrid Bergmann in case Ref. 7 O 234/26. The next hearing will take place on 04/22/2026 at the Frankfurt Regional Court.

What the AI sees:

[PERSON_1] (**[EMAIL_1]**) represents **[PERSON_2]** in case **[CASE_NO_1]**. The next hearing will take place on **[DATE_1]** at the **[ORG_1]**.

The AI still understands the context perfectly. It knows that this is about a court hearing, who represents whom, and when the hearing is scheduled. It just no longer knows *who* exactly is involved.

Government, Public Administration & Public Sector



Ministries, state agencies, municipalities, and subordinate authorities process citizen data at the highest protection level. At the same time, Germany's Online Access Act (OZG) mandates digitalization. Several German states have already adopted AI strategies, yet data protection officers block every direct API connection to US AI services. A political dilemma that can be solved technically.

The Challenge

- Citizen data is subject to BSI IT baseline protection and VS-NfD (restricted classified information) requirements
- Registration data, social benefit applications, building permits, and tax data must not reach third parties
- Police and security authorities process investigation data and witness protection information
- Educational institutions process student data and exam results
- State data protection laws prohibit cloud AI involving personal data

Anonix in Practice

PROMPT:
Rewrite this administrative notice in plain, citizen-friendly language.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Dear <u>Ms. Ingrid Bergmann</u>, pursuant to Sections 1 and 3 WoGG (Housing Benefit Act), your application (Ref. <u>WG-2026-04817</u>) is granted. The benefit will be transferred to account <u>DE89 3704 0044 0532 0130 00</u>.	Dear <u>[PERSON_1]</u>, pursuant to Sections 1 and 3 WoGG (Housing Benefit Act), your application (Ref. <u>[CASE_NO_1]</u>) is granted. The benefit will be transferred to account <u>[IBAN_1]</u>.	Dear <u>Ms. Ingrid Bergmann</u>, Good news: your housing benefit application (<u>WG-2026-04817</u>) has been approved. We will transfer the money directly to your account <u>DE89 3704 0044 0532 0130 00</u>. You don't need to do anything else.

Hospitals and Clinics



Under Art. 9 GDPR, health data enjoys the highest level of protection. Findings, diagnoses, medication plans, and treatment histories are special categories of personal data. A violation isn't just a fine—it destroys the trust between doctor and patient. At the same time, AI-powered systems promise enormous advances in diagnostics, documentation, and treatment planning.

The Challenge

- Art. 9 GDPR: health data is a special category whose processing is prohibited in principle
- German state hospital laws impose additional requirements
- Hospital information systems (HIS) contain complete patient records
- Medical confidentiality under Section 203 of the German Criminal Code (StGB): disclosure to third parties is a criminal offense
- Enormous documentation burden: physicians spend up to 40% of their working time on documentation

Anonix in Practice

PROMPT:
Turn these notes into a structured physician's letter to the family doctor.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Patient <u>Thomas Weber</u> , DOB <u>03/14/1968</u> . Diagnosis: <u>type 2 diabetes</u> . HbA1c: <u>8.2%</u> . Medication: <u>metformin 1000 mg</u> .	Patient <u>[PERSON_1]</u> , DOB <u>[DATE_1]</u> . Diagnosis: <u>[DIAGNOSIS_1]</u> . HbA1c: <u>[VALUE_1]</u> . Medication: <u>[MEDICATION_1]</u> .	Dear Colleague, We are writing regarding our patient <u>Thomas Weber</u> , DOB <u>03/14/1968</u> . Diagnosis: <u>type 2 diabetes</u> Labs: HbA1c <u>8.2%</u> Medication: <u>metformin 1000 mg</u> Plan: We recommend adjusting the therapy and a follow-up visit in 3 months.

Care Facilities



Germany has more than 15,000 nursing homes and home care services. The skilled labor shortage is acute: over 35,000 positions are unfilled. Caregivers spend up to 30% of their working time on documentation instead of with residents. AI can reduce this burden—but care documentation contains health data under Art. 9 GDPR that must not reach external services.

The Challenge

- Care documentation, medication plans, and care logs contain Art. 9 data
- Resident data includes diagnoses, cognitive impairments, and behavioral issues
- Inspections by the Medical Service (MDK) require complete, audit-proof documentation
- Care home supervisory authorities and data protection authorities regularly review data processing

Anonix in Practice

PROMPT:
Turn these notes into a structured handover report for the night shift.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Resident <u>Ms. Helga Schneider</u> , room <u>214</u> , increased fall risk. Medication <u>Marcumar</u> adjusted to <u>2.5 mg</u> .	Resident <u>[PERSON_1]</u> , room <u>[ROOM_1]</u> , increased fall risk. Medication <u>[MEDICATION_1]</u> adjusted to <u>[DOSAGE_1]</u> .	Night Shift Handover <u>Ms. Helga Schneider</u> (Rm. <u>214</u>): - Increased fall risk—raise bed rails - <u>Marcumar</u> adjusted to <u>2.5 mg</u> - Close monitoring required - Schedule next INR check

Courts and Justice



The German justice system handles millions of proceedings every year. Case files, judgments, indictments, and witness statements contain the most sensitive personal data there is: victim data, witness protection information, statements by minors. Judicial confidentiality prohibits any uncontrolled disclosure. At the same time, the pressure is growing to shorten proceedings through AI-assisted research.

The Challenge

- Case files contain the names of defendants, victims, witnesses, and expert witnesses
- Juvenile criminal law: special protection for underage defendants and victims
- Witness protection programs: any disclosure of data can endanger lives
- Judicial independence requires that no external third parties gain access to case data

Anonix in Practice

PROMPT:
Draft the operative part of a court order for these facts.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Defendant <u>Marcus Klein</u> , Ref. <u>7 O 234/26</u> . Charge: fraud under Section 263 StGB. Injured party: <u>Müller GmbH, Frankfurt</u> .	Defendant <u>[PERSON_1]</u> , Ref. <u>[CASE_NO_1]</u> . Charge: fraud under Section 263 StGB. Injured party: <u>[ORG_1]</u> , <u>[LOCATION_1]</u> .	Order In the criminal case against <u>Marcus Klein</u> , Ref. <u>7 O 234/26</u> , for fraud under Section 263 StGB to the detriment of <u>Müller GmbH</u> , a date for the main hearing is set at the Regional Court of <u>Frankfurt</u> .

Social Services and Youth Welfare Offices



Youth welfare offices and social service agencies process the most sensitive data in our society: child endangerment, taking children into care, addiction issues, domestic violence. Case reports and assessments concern people in crisis. A data breach here is not an abstract compliance violation—it can destroy families and put children at risk.

The Challenge

- Case reports contain highly sensitive data about children, parents, and family situations
- Section 65 of Book VIII of the German Social Code (SGB VIII): social data is subject to heightened confidentiality protection
- Assessments contain psychological evaluations and psychiatric diagnoses
- Caseworkers are chronically overloaded; AI could ease the documentation burden

Anonix in Practice

PROMPT:
Turn these notes into a structured case summary with recommended measures under SGB VIII.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
<u>Yilmaz</u> family, <u>3 children</u> (ages <u>4, 7, 12</u>). Report by school, suspected child endangerment. Home visit on <u>02/12/2026</u>.	<u>[PERSON_1]</u> family, <u>[COUNT_1]</u> children (ages <u>[AGE_1]</u>). Report by school, suspected child endangerment. Home visit on <u>[DATE_1]</u>.	Case Summary Family: <u>Yilmaz</u> Children: <u>3</u> (ages <u>4, 7, 12</u>) Report: school Suspicion: Section 8a SGB VIII Home visit: <u>02/12/2026</u> Recommendation: Initiate parenting counseling under Section 28 SGB VIII

Law Firms



Attorney-client confidentiality is not only protected by professional rules—a breach is a criminal offense under Section 203 of the German Criminal Code (StGB). Draft contracts, briefs, due diligence reports, and client data must never be disclosed to third parties under any circumstances. At the same time, AI offers enormous potential: contract analysis, research, brief review, and legal tech are fundamentally changing the industry. Law firms that don't use AI lose clients.

The Challenge

- Section 203 StGB: disclosing data is a criminal breach of attorney-client confidentiality
- Contracts contain company names, trade secrets, purchase prices, and terms
- M&A transactions: due diligence documents of the highest confidentiality
- In the event of a data breach, opposing counsel could claim a conflict of interest

Anonix in Practice

PROMPT:
Draft a liability clause based on this contract data.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Seller: <u>TechVision GmbH</u> , represented by <u>Dr. Stefan Hofmann</u> . Purchase price: <u>EUR 4.2 million</u> . Liability cap: <u>EUR 840,000</u> .	Seller: <u>[ORG_1]</u> , represented by <u>[PERSON_1]</u> . Purchase price: <u>[AMOUNT_1]</u> . Liability cap: <u>[AMOUNT_2]</u> .	Section 7 Liability (1) The liability of <u>TechVision GmbH</u> , represented by <u>Dr. Stefan Hofmann</u> , is limited to <u>EUR 840,000</u> . (2) The buyer's claims become time-barred 24 months after handover. (3) This limitation does not apply in cases of intent or gross negligence.

Tax Advisors and Auditors



Tax confidentiality under Section 203 StGB protects client data under criminal law. Tax returns, balance sheets, salary information, and tax audit documents are among a company's most confidential business data. A violation doesn't just threaten the advisor's license—it can trigger liability claims that threaten their very existence.

The Challenge

- Tax confidentiality under Section 203 StGB and Section 57 of the German Tax Advisory Act (StBerG)
- Client data includes revenue, salaries, assets, and tax IDs
- Tax audit documents contain detailed business figures
- The German Public Accountant Act (WPO) and ISA standards require strict information security for auditors

Anonix in Practice

PROMPT:
Draft an objection letter to the tax office for this client.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Client <u>Bäckerei Kraus GmbH</u> , tax no. <u>013/123/45678</u> . Objection to 2025 VAT assessment. Vehicle: 60% business, 40% private.	Client <u>[ORG_1]</u> , tax no. <u>[TAX_ID_1]</u> . Objection to 2025 VAT assessment. Vehicle: 60% business, 40% private.	Objection Tax Office ... Tax no.: <u>013/123/45678</u> On behalf of and as authorized by our client <u>Bäckerei Kraus GmbH</u>, we hereby object to the 2025 VAT assessment. Grounds: With 60% business use, a proportional input tax deduction under Section 15 UStG must be recognized. We request a suspension of enforcement.

Banks and Insurance Companies



BaFin-regulated companies are subject to banking secrecy and, since January 2025, under DORA (Digital Operational Resilience Act), to the strictest IT security requirements in the EU. Account data, loan applications, claims reports, and risk assessments require the highest level of protection. At the same time, banks and insurers are investing billions in AI-powered process automation.

The Challenge

- Banking secrecy under Section 383 of the German Code of Civil Procedure (ZPO) and the banks' general terms and conditions
- DORA: strict requirements for IT risk management and third-party oversight
- MaRisk/BAIT: outsourcing rules for IT services
- Customer KYC data, account transactions, and credit scoring results are highly sensitive

Anonix in Practice

PROMPT:

Turn this claims report into a settlement letter to the policyholder.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Policyholder <u>Sabine Hartmann</u> , policy no. <u>KFZ-2026-881234</u> . Accident on <u>03/03/2026</u> in <u>Munich, Leopoldstr. 42</u> . Damage: <u>EUR 12,400</u> .	Policyholder <u>[PERSON_1]</u> , policy no. <u>[POLICY_NO_1]</u> . Accident on <u>[DATE_1]</u> in <u>[LOCATION_1]</u> , <u>[ADDRESS_1]</u> . Damage: <u>[AMOUNT_1]</u> .	Dear <u>Ms. Hartmann</u> , Regarding your claim of <u>03/03/2026</u> (policy no. <u>KFZ-2026-881234</u>), we are pleased to inform you that the settlement process has been initiated. Location of loss: <u>Munich, Leopoldstr. 42</u> Reported damage: <u>EUR 12,400</u> An adjuster will contact you within 5 business days.

HR Departments



Every company—regardless of industry—processes highly sensitive employee data in its HR department. Personnel files, applicant data, salary information, sick notes, and employee evaluations are personal data whose disclosure to third parties has consequences under both employment and data protection law. AI-powered recruiting, performance reviews, and workforce planning are booming—often without the data protection officer’s knowledge.

The Challenge

- Section 26 of the German Federal Data Protection Act (BDSG): strict purpose limitation when processing employee data
- Applicant data: compliance with the German General Equal Treatment Act (AGG) and discrimination risks in AI-assisted pre-screening
- Works agreements often explicitly regulate the use of AI tools
- Salary data, sick notes, and formal warnings are particularly sensitive

Anonix in Practice

PROMPT:
Turn this internal job description into an external job posting.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Successor for <u>Martin Schäfer</u> , Controlling dept. Salary: <u>EUR 78,000</u> . Manager: <u>Dr. Lisa Berger</u> . Resignation effective <u>06/30/2026</u> .	Successor for <u>[PERSON_1]</u> , Controlling dept. Salary: <u>[AMOUNT_1]</u> . Manager: <u>[PERSON_2]</u> . Resignation effective <u>[DATE_1]</u> .	Controller (m/f/d) Your responsibilities: - Budget planning and forecasting - Monthly and annual closings - Reporting to <u>Dr. Lisa Berger</u> What we offer: - Compensation: <u>EUR 78,000</u> - Start: as soon as possible

Educational Institutions



Universities, colleges, schools, and research institutions are under double pressure: on the one hand, they are expected to be AI pioneers and provide students with modern tools. On the other hand, they process student data, exam results, and personal research data covered by state data protection laws. Several state data protection commissioners have already published recommendations against using US cloud AI services.

The Challenge

- State data protection laws and higher education acts govern the handling of student data
- Exam results, student ID numbers, and reasons for withdrawal are specially protected
- Research data involving personal information is subject to ethics committees
- Lecturers are already using AI extensively—often without institutional approval

Anonix in Practice

PROMPT:
Turn this course evaluation into a structured improvement report for the academic committee.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Course <u>BWL-301</u> , Prof. <u>Dr. Andreas Keller</u> . Grade: <u>2.3</u> . Comment: 'The lecturer explains well, but <u>Ms. Schneider</u> from the tutorial was unprepared.'	Course <u>[COURSE_1]</u> , Prof. <u>[PERSON_1]</u> . Grade: <u>[VALUE_1]</u> . Comment: 'The lecturer explains well, but <u>[PERSON_2]</u> from the tutorial was unprepared.'	Evaluation Report Course: <u>BWL-301</u> Lecturer: Prof. <u>Dr. Andreas Keller</u> Overall grade: <u>2.3</u> Strengths: Clear explanations in the lecture. Action needed: Involve and prepare the tutorial leader (<u>Ms. Schneider</u>) more closely.

Pharma and Life Sciences



The pharmaceutical and life sciences industries process clinical trial data, patient information, and proprietary research results. GxP regulations, FDA 21 CFR Part 11, and the EU Clinical Trials Regulation set the highest standards for data integrity and confidentiality. AI speeds up drug development by years—but trial data must not fall into the wrong hands.

The Challenge

- Clinical trial data contains patient names, diagnoses, and treatment histories
- Proprietary compound data and research results are trade secrets
- GxP compliance requires complete traceability of all data processing
- Adverse event reports contain personal reports of side effects

Anonix in Practice

PROMPT:

Turn this data into a structured adverse event report in CIOMS format.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Subject <u>ID-4821</u> , investigator: <u>Dr. Maria Engel</u> . Adverse reaction: <u>grade 2 nausea</u> after <u>Compound XR-7</u> .	Subject <u>[ID_1]</u> , investigator: <u>[PERSON_1]</u> . Adverse reaction: <u>[SYMPTOM_1]</u> after <u>[COMPOUND_1]</u> .	CIOMS Adverse Event Report Subject: <u>ID-4821</u> Investigator: <u>Dr. Maria Engel</u> Event: <u>grade 2 nausea</u> Suspect drug: <u>Compound XR-7</u> Causality: possible SAE: No Action: symptomatic treatment, continued monitoring

Software Companies and IT Service Providers



Software companies and IT service providers integrate LLM APIs into their own products—for customers in every industry. They are responsible for ensuring that their end customers’ data remains protected. A data protection violation affects not only their own company but every customer whose data flows through the API. Here, Anonix becomes an infrastructure component built into the product itself.

The Challenge

- Data processing under Art. 28 GDPR: the software provider is liable for its customers’ data
- Customer data from regulated industries (healthcare, finance, legal) flows through the API
- SOC 2 and ISO 27001 certifications require demonstrable data protection measures
- Customers increasingly demand proof of GDPR-compliant AI use in tenders

Anonix in Practice

PROMPT:
Use this application data to write a structured acknowledgment of receipt to the applicant.

ORIGINAL DATA	SENT TO THE AI	RESULT (DE-ANONYMIZED)
Application from <u>Anna Fischer</u> , <u>anna.fischer@email.de</u> , for <u>Senior Controller</u> at <u>Autohaus Schmidt GmbH</u> . Desired salary: <u>EUR 85,000</u> .	Application from <u>[PERSON_1]</u> , <u>[EMAIL_1]</u> , for <u>[POSITION_1]</u> at <u>[ORG_1]</u> . Desired salary: <u>[AMOUNT_1]</u> .	Dear <u>Ms. Fischer</u> , Thank you for your application for the position of <u>Senior Controller</u> at <u>Autohaus Schmidt GmbH</u> . We have received your documents and are reviewing them carefully. You will hear from us within 10 business days. Sincerely, The HR Team at <u>Autohaus Schmidt GmbH</u>

Why Existing Solutions Fall Short



There are certainly products that can detect personal data. But none of them solves the actual problem: using AI APIs securely without data leaving the organization.

CRITERION	ANONIX	OPENROUTER®	MICROSOFT® PRESIDIO®	GOOGLE® CLOUD DLP	AWS® COMPREHEND
Transparent AI proxy	Yes	Yes (AI gateway)	No	No	No
Supported LLM providers	OpenAI® + all compatible, Claude®, Gemini®	400+ models	No proxy	No proxy	No proxy
Data stays in the EU	Yes (EU SaaS)	EU routing from Business plan	Yes	Google® Cloud	AWS® Cloud
Automatic de-anonymization	Yes	No, redaction only	No	No	No
Industry-specific detection	Zero-Shot NER	8 types + regex	Standard only	150+ types	Standard only
Multilingual (> 5 languages)	8 languages	Not documented	Partial	Yes	2 languages
AI-powered quality check	SLM review	No	No	No	No
Real-time streaming	SSE support	No de-anonymization	No	No	No
Typo tolerance	Fuzzy matching	No	No	No	No
EU-compliant deployment	EU SaaS	US Cloud Act	Open source	US Cloud Act	US Cloud Act

Where Is the Difference?

OpenRouter® is an AI gateway with over 400 models and has recently added redaction of sensitive data. But the check only happens on the servers of the US company, so the data leaves the organization in plain text. The redaction is permanent: every name becomes [PERSON_NAME], the AI can no longer tell people apart, and the response contains no original data. OpenRouter® detects names and addresses only as a beta feature. If this check times out, the request is forwarded unredacted.

Google® Cloud DLP and **AWS® Comprehend** do detect personal data, but they are themselves cloud services of US companies. Taking a detour through AWS® Bedrock or Microsoft® Azure® in an EU data center doesn't change anything either: the CLOUD Act applies regardless of server location, because the parent company is based in the US.


Microsoft® Presidio® is open source and can be run locally. But it has no proxy function, no de-anonymization, and no zero-shot detection. It is a toolkit, not a finished product.

Anonix is the only system that combines all three requirements in one product: privacy-compliant AI detection, transparent LLM proxy integration, and complete de-anonymization of the AI response.

The Architecture at a Glance



Anonix consists of five components that are operated as a fully managed SaaS service in the EU. Not a single byte of personal data leaves the European infrastructure.



The Five Building Blocks

COMPONENT	FUNCTION	TECHNOLOGY
Anonix Backend	Proxy engine, anonymization, user management, API	Python® / FastAPI
Anonix Frontend	Admin interface for providers, rules, audit log	React
Anonix ConvNet NER	Neural network for detecting names, places, companies	CNN microservice (8 languages)
Anonix Zero-Shot NER	Industry-specific detection without training	Transformer microservice
AI Quality Check	Final review layer for overlooked data	Small language model (local or API)

Security at Every Level

- AES-256-GCM encryption of all stored API keys
- Argon2id password hashing (military-grade standard)
- Isolated databases per tenant
- Two-factor authentication (TOTP)
- Rate limiting and CSRF protection
- Structured audit logging

Investment and Cost-Effectiveness



Data protection doesn't have to be a cost driver. Anonix costs a fraction of what a single data protection violation would cost. And at the same time, it enables the productive use of AI that your employees are already asking for.

What a Violation Costs

530M

EUR FINE

TikTok (2025) for unlawful data transfers

45M

EUR FINE

Highest penalty in Germany (Vodafone, 2025)

4%

OF ANNUAL REVENUE

Maximum GDPR penalty for companies

What Anonix Costs

Anonix charges by usage—not by users. A base fee plus a price per anonymization request. Unlimited users in every plan.

VOLUME / MONTH	PRICE PER REQUEST	IN EUROS	NOTES
Up to 50,000	1.4 cents	EUR 0.014	Entry level / small customers
50,000 – 250,000	0.7 cents	EUR 0.007	Standard / mid-sized companies
Over 250,000	0.5 cents	EUR 0.005	Enterprise / large customers

Base fee: EUR 99 – 999 / month depending on plan and company size. All prices plus VAT. SLM premium module (AI quality check with dedicated GPU inference): +1.4 cents per request.

Example: 100 Employees

Employees	100
Requests/month	22,000
Base fee	EUR 99
22,000 × EUR 0.014	EUR 308
Total / month	EUR 407

Example: 500 Employees

Employees	500
Requests/month	110,000
Base fee	EUR 199
110,000 × EUR 0.007	EUR 770
Total / month	EUR 969

Example: 1,000 Employees

Employees	1,000
Requests/month	220,000
Base fee	EUR 499
220,000 × EUR 0.007	EUR 1,540
Total / month	EUR 2,039

Example: 5,000 Employees

Employees	5,000
Requests/month	1,100,000
Base fee	EUR 999
1,100,000 × EUR 0.005	EUR 5,500
Total / month	EUR 6,499

Basis: ~10 AI requests per employee per working day × 22 working days = approx. 220 requests per employee per month. The more your company uses AI, the cheaper each individual request becomes.

All prices are net prices plus VAT.

The math is simple: The annual cost of Anonix is a fraction of a single GDPR fine. And it costs less than a single day of lost productivity if you have to ban AI entirely.

Compliance at the Push of a Button



Data protection is not a project with a beginning and an end. It is an ongoing process. Anonix automates the technical measures that regulators require.

GDPR Compliance

- ✓ **Art. 5(1)(c) (data minimization)**—Only placeholders reach the AI provider. Personal data is reduced to the technically necessary minimum.
- ✓ **Art. 25 (data protection by design)**—Anonymization takes place before the API call. Without the proxy configuration, unprotected use is not possible.
- ✓ **Art. 32 (security of processing)**—AES-256 encryption, Argon2id hashing, two-factor authentication, and isolated tenant databases.
- ✓ **Art. 28 (processors)**—The AI provider sees no personal data. The requirements placed on the processor are significantly reduced.
- ✓ **Art. 30 (records of processing activities)**—Complete audit log with full logging of all proxy requests.

CLOUD Act and Third-Country Transfers

Architectural guarantee: Even if US authorities compel the AI provider to hand over data, they receive only anonymized text fragments with placeholders. Anonix stores no mapping between placeholders and original data. The mapping exists only in memory during processing and is discarded immediately afterward. The original data remains exclusively with the customer. Re-identification is technically impossible.

EU AI Act (Regulation 2024/1689)

Anonix is a data minimization tool and does not fall into any high-risk category of the EU AI Act. It makes no automated decisions about individuals, creates no profiles, and does not affect anyone's rights. Rather, it reduces the risks posed by the actual high-risk applications (the external AI models).

Industry-Specific Compliance

REGULATION	INDUSTRY	HOW ANONIX HELPS
Section 203 StGB	Lawyers, tax advisors, physicians	Professional secrecy is preserved, since no client/patient data is transmitted
DORA (EU 2022/2554)	Banks, insurance companies	ICT risk management through controlled AI use
State data protection laws	Public administration	Citizen data stays on municipal/state infrastructure
SGB (social data protection)	Social services, youth welfare offices	Social data is anonymized before external processing

Voices from the Field



The challenge is the same across industries: how can organizations capture the productivity gains of AI without losing control over sensitive data? Three perspectives.

"We had two options: ban AI entirely or find a way to use it in compliance with data protection law. The first would have set us back against the competition. With Anonix, our employees can now officially use AI tools. Our data protection officer reviewed and approved the concept. The deciding factor was that personal data truly never leaves our servers."

DR. KATHARINA SOMMER, DATA PROTECTION OFFICER, MID-SIZED INSURANCE COMPANY

"As head of IT for a municipality, I had a clear problem: our business applications were supposed to get AI-powered text processing, but the data protection officer blocked every direct API connection to US providers. Anonix gave us a clean solution. We switched the API keys in our systems over to the Anonix proxy, and that was it. The SaaS service runs in the EU, and not a single citizen's name goes to the AI providers anymore."

MICHAEL BRANDT, HEAD OF IT, CITY ADMINISTRATION (APPROX. 120,000 RESIDENTS)

"Our firm has 14 attorneys. Every one of us could use AI to review briefs, summarize rulings, or draft contract clauses. But attorney-client confidentiality comes first. With Anonix, I can now tell my colleagues with a clear conscience: use the AI. Your client data is protected. And the best part: it took less than two hours to get everything up and running."

DR. FLORIAN HARTMANN, PARTNER, BUSINESS LAW FIRM

Note: The testimonials presented here illustrate typical use cases and challenges as described in conversations with potential users from the respective industries.

Your Path to Anonix



Anonix is designed to take you from first contact to productive use in just a few days. No lengthy implementation projects, no rebuilding of your existing systems.

1

Consultation
We analyze your requirements:
Which AI providers do you use?
Which data types need to be protected? Which industry-specific rules apply?

2

Pilot Phase
Access to your Anonix SaaS account. Configuration of the anonymization rules. Trial operation with one team.

3

Production
Rollout to all workstations. Swap the API keys, done. Ongoing support, automatic updates, and assistance included.

Deployment

Anonix is provided as a fully managed SaaS service. Hosting, updates, maintenance, and support are included in the package. You can focus on your core business while we run the infrastructure—exclusively on EU servers, of course.

OPTION	DESCRIPTION	SUITABLE FOR
Anonix SaaS	Fully managed solution on EU infrastructure. Hosting, updates, and support included. Ready to use immediately.	Companies, public authorities, law firms, medical practices
Enterprise On-Premises	For organizations with special data sovereignty requirements, we also offer a dedicated on-premises installation on request.	On request

Ready for the Next Step?

Let's find out together how Anonix fits into your organization.

info@anonix.ai • www.anonix.ai

Legal Notice



Copyright, Trademark, and License Rights

All copyright, trademark, and license rights to Anonix Privacy Proxy and this document are held by:

Company	Convecto GmbH
Address	Ludwigstraße 180d, 63067 Offenbach, Germany
Phone	+49 69 40897270
Web	www.convecto.com

Trademarks

Anonix, Anonix Privacy Proxy, Anonix ConvNet NER, and Anonix Zero-Shot NER are trademarks of Convecto GmbH. OpenAI®, GPT®, Anthropic®, Claude®, Google®, Gemini®, Microsoft®, DeepSeek®, Mistral®, xAI®, Grok®, Groq®, Perplexity®, Together AI®, OpenRouter®, Ollama®, LM Studio®, Cursor®, GitHub® Copilot®, LangChain®, LlamaIndex®, and AWS® are trademarks of their respective owners.

© 2000–2026 Convecto GmbH. All rights reserved.

The information contained in this document has been compiled with the utmost care. However, no guarantee is given as to its accuracy, completeness, or timeliness.