

ACADEX

Technical and Organisational Measures

pursuant to Article 32 GDPR

Version 1.0

Status: 28 September 2026

Working draft

Acadex Solutions UG (haftungsbeschränkt)

Reuterweg 5 | 49637 Menslage | Germany

Note on this working draft

This version reflects the technical state intended for publication. Measures that are to be implemented before release are described as existing measures only where their implementation before publication has been firmly scheduled. The server-side website statistics and Product Metrics described in this version are treated as part of the implemented release state. The document should undergo legal review before contractual use.

1. Purpose and scope

This document describes the technical and organisational measures used by Acadex to protect personal data processed through the Acadex platform.

The measures are intended in particular to ensure a level of security appropriate to the relevant risk with regard to the confidentiality, integrity, availability and resilience of the systems and services used for processing.

The measures described reflect the current technical and organisational state of the service. Measures that are merely planned or not yet fully implemented are not presented as existing safeguards.

The specific processing may vary depending on the configuration of the relevant account or institution. This applies in particular to the selection of external model providers and the associated processing locations.

2. Confidentiality

2.1 Administrative access

Administrative access to production systems is limited to the technical lead and, where applicable, other persons who require such access for their respective duties.

Administrative server access is performed exclusively via SSH using individual cryptographic keys. Password-based SSH login is disabled.

Administrative user accounts are assigned individually and are not shared between multiple persons. Permissions are granted only in accordance with operational need.

The administrative accounts used for hosting, server administration and source-code management at STRATO, Plesk and GitHub are protected by multi-factor authentication.

Administrative activities and production deployments are currently carried out by the technical lead.

2.2 User authentication

User accounts are currently authenticated using an email address and password.

Passwords are not stored in plaintext and are hashed exclusively using Argon2id.

Short-lived, cryptographically signed access tokens are used for authenticated access to the application.

Personal accounts are assigned to individual persons. Conversations, memories, watchlists, documents and other personal content are associated with the relevant account.

2.3 Encryption in transit

Connections between clients and Acadex, and between Acadex and connected processors, use encrypted transport connections based on TLS.

Unencrypted transmission of personal application data is not intended during normal operation.

2.4 Encryption of stored content

The following content data is encrypted within the application database using AES-256-GCM:

- conversation content
- saved memories
- content of uploaded documents
- chat titles
- request content stored for individual processing operations

A separate nonce is used for each encryption operation.

The key required for decryption is stored separately from the application database.

An isolated copy of the application database therefore does not contain all information required to decrypt this content.

2.5 Management of secrets and credentials

Credentials, API keys and cryptographic secrets required on an ongoing basis are not stored in the source code.

They are stored in a dedicated encrypted password manager. The key material required to access that password manager is kept on a separate physical medium, with a secured backup copy of that key material.

For ongoing operation, required secrets are supplied at runtime by the technical lead as environment variables.

Access to these secrets is currently limited to the technical lead.

Secrets and credentials are rotated approximately every three months and additionally when circumstances require it, in particular where loss of confidentiality is suspected or identified.

2.6 Data not encrypted in the same manner

Not all stored information can be protected at application level using the same form of encryption.

In particular, the text used for semantic search within saved memories must remain machine-searchable and is therefore currently not stored using the same field-level AES-256-GCM encryption.

Public bibliographic records, technical identifiers, timestamps and operational metadata required for the service are likewise not protected by this field-level content encryption.

3. Separation and access control

3.1 Separation of personal content by account

Personal content is assigned to a specific user account.

Database queries accessing user-specific content are restricted to the relevant account.

Personal accounts within an institutional account do not automatically share conversations, memories or watchlists with each other.

3.2 Separation of system environments

Production, development and test environments use separate databases.

Production data is not routinely used to populate development or test environments.

3.3 Network separation

The application, database, cache, search index and internal inference services operate in separate isolated containers within a private network.

Only the designated access layer or reverse proxy is exposed externally.

Internal services are not directly exposed to the public network.

4. Limitation of external processing

4.1 Data minimisation for external model providers

Where an external model provider is used for a request, only information required for the relevant processing is transmitted.

This may include in particular:

- the current user input
- relevant conversation context
- memories retrieved as relevant to the request
- relevant bibliographic information
- document content required to perform the task

The complete data held for a user account is not made available to a model provider by default.

4.2 Permitted model providers

Acadex supports models or APIs from Mistral, OpenAI, Anthropic and Google, either directly or through controlled cloud processing routes.

The account configuration determines which model providers and processing profiles may be used for a particular account.

For institutional accounts, this allows requirements concerning processing location or permitted subprocessors to be enforced technically.

Technical support for a model provider within Acadex does not automatically mean that the provider processes personal data for every customer.

4.3 Locally executed processing

Embeddings are generated on infrastructure controlled by Acadex.

The relevant content therefore does not need to be transmitted to an external embedding provider for this purpose.

Other locally executed services are documented according to the processing they actually perform.

5. Integrity, logging and usage metrics

5.1 Protection against unintended changes

Access to personal content takes place through authenticated application and administrative interfaces.

Database access by the application is performed within the applicable account and authorisation context.

5.2 Logging and error analysis

Conversation content is not written to operational logs.

Internal user identifiers may be logged for technical error analysis so that an error can be associated with a specific account or processing operation.

IP addresses are processed during operation to enforce technical safeguards, in particular rate limits. They are not used for advertising purposes or user profiling.

Operational logs remain on the relevant production system for seven days.

Required logs are then compressed and encrypted and transferred to a separate archive server. Archived logs are retained for no longer than three months and are then deleted.

Access to operational and archived logs is restricted to administrative accounts.

5.3 Website statistics

Acadex generates aggregated statistics on views of selected website areas on its own infrastructure. These statistics are used for basic reach and product management and are generated without an external analytics provider.

In particular, the analytics dataset does not store persistent visitor identifiers, browser fingerprints, URL query parameters, search terms or complete referrer URLs. Full IP addresses are not retained as persistent analytics attributes.

The statistics primarily count defined pages or page categories and time periods. Website usage statistics are technically and organisationally separated from operational server logs.

5.4 Product Metrics

To assess basic usage and business development, Acadex processes a limited set of clearly defined product and contractual states.

These include in particular registration, first successful use of a core product function (activation), the start of or change to a plan, and cancellation. Aggregated metrics such as activation rate, conversion, retention, churn, MRR or ARR may be derived from these states.

Where possible, these states are derived from existing account, contract, billing or operational data. Acadex does not operate an additional comprehensive event-tracking infrastructure or a complete clickstream for this purpose.

5.5 Exclusion of content data

In particular, the following data is not analysed for Product Analytics:

- conversation content and user inputs
- generated answers
- content of uploaded documents
- saved memories
- research topics and bibliographic search terms
- individual click or navigation paths

This separation is implemented through the selection of recorded states and the structure of the metrics data.

5.6 Deletion and aggregation of Product Metrics

Account-linked Product Metrics data is deleted together with the relevant account unless a statutory retention obligation or another independent processing purpose requires otherwise.

Statistics already generated may be retained where they have been aggregated so that they can no longer be associated with individual persons.

Long-term statistical analysis is generally based on such aggregated metrics.

5.7 Data minimisation for new metrics

When new metrics are introduced, Acadex assesses whether the intended insight can be obtained without collecting additional personal event data. Where a metric can be derived from existing account, contract, billing or operational data, that approach is preferred over introducing additional tracking events.

6. Deletion and data lifecycle

6.1 Primary and secondary data

PostgreSQL serves as the primary data store for application and user data that cannot be reconstructed from other systems.

Qdrant and Elasticsearch contain derived or indexed data. The information stored there can be regenerated from the corresponding primary data.

Bibliographic research data can likewise be reconstructed from the external data sources used by Acadex.

This distinction is taken into account for backup, restoration and deletion.

6.2 Deletion by users

Users can delete individual conversations, memories and uploaded documents within the application without having to contact support.

6.3 Account deletion

When an account is deleted, the user data associated with that account, including account-linked Product Metrics states, is removed from the active application systems unless a statutory retention obligation or another independent processing purpose requires otherwise.

Referential deletion within the relational data store is implemented through database relationships and cascading deletion rules and does not depend solely on application code addressing each affected table individually.

Derived user data in Qdrant and Elasticsearch is also removed as part of the deletion process or regenerated from the remaining primary data.

6.4 Verification of deletion logic

The database schema is checked automatically to determine whether tables containing user-linked data are covered by the intended deletion logic.

If such a table is detected without the required integration into the deletion logic, the check fails.

6.5 Deletion from backups

Data already deleted from active systems may remain in existing backup copies until the relevant retention period expires.

Backups are not used for normal production operation. At the end of their intended retention period, they are deleted or replaced as part of backup rotation.

Restoring an older backup does not result in data that was validly deleted in the meantime being permanently reintroduced into normal production use.

7. Availability, backup and recovery

7.1 Hosting

Acadex production systems are operated on infrastructure provided by STRATO GmbH in Germany.

The application, database and backups held by Acadex are operated or stored on infrastructure within Germany.

7.2 Current system architecture

The production application currently runs on a single application instance.

The architecture is technically capable of supporting multiple instances. A production high-availability configuration has not yet been implemented.

Failure of this instance or of central infrastructure components can therefore currently interrupt the service.

Acadex does not currently commit to guaranteed availability or a general Service Level Agreement.

7.3 Backups by data class

Acadex distinguishes between primary user data and reconstructable secondary or research data for backup purposes.

User data that cannot readily be reconstructed is stored primarily in PostgreSQL and backed up independently of the complete research-data environment.

Automated PostgreSQL backups are created at a significantly shorter interval than full backups of the complete data environment.

PostgreSQL backup interval: [TO BE COMPLETED]

PostgreSQL backup retention period: [TO BE COMPLETED]

Full backups of the complete data environment are created approximately every three to four months.

The complete data environment exceeds 3 TB and consists largely of bibliographic data, search indexes, embeddings and other information that can be reconstructed from primary sources or the primary application data.

Qdrant and Elasticsearch data, as well as bibliographic data, can therefore be rebuilt or restored from the underlying sources where necessary.

7.4 Recovery

Recoverability of PostgreSQL backups is considered separately from restoration of the complete research-data environment.

Procedure and interval for PostgreSQL restore tests: [TO BE COMPLETED]

Acadex does not currently commit to binding Recovery Time Objectives (RTO) or Recovery Point Objectives (RPO) for restoration of the complete data environment.

7.5 Monitoring

Production systems are currently monitored administratively.

Automated external uptime monitoring with alerting is not currently in operation and is therefore not described as an implemented measure in this document.

Additional structured monitoring and alerting are planned. Concrete implementation and operational scope: [TO BE COMPLETED]

8. Development and deployment

Acadex source code is version-controlled through GitHub and maintained in a structured monorepository.

Changes are developed in separate branches and incorporated into the designated main branches through pull requests. Code review forms part of the development process.

The application uses automated unit tests and User Acceptance Tests for selected functionality.

A CI/CD pipeline supports build, testing and deployment.

Production deployments are currently carried out by the technical lead.

The application is operated in a containerised environment based on Docker.

Operating-system updates are managed through the Plesk server-administration environment.

Updates to Docker images and the software components they contain are currently performed on an event-driven basis.

9. Subprocessors

External service providers that process personal data on behalf of Acadex as part of their services are engaged in accordance with applicable data-protection requirements.

The subprocessors used, their functions and the relevant processing locations are documented separately.

Which model providers may actually receive personal content from a particular account depends on that account's configuration.

Support for a provider within the Acadex platform does not automatically mean that the provider processes data for every customer.

Changes concerning subprocessors are handled in accordance with the Data Processing Agreement.

10. Security incident management

Operational responsibility for technical security incidents lies with the technical lead.

External communication with affected customers or institutions is coordinated with the person responsible for customer communication.

Where Acadex acts as a processor and becomes aware of a personal data breach affecting data processed on behalf of a controller, Acadex notifies the relevant controller without undue delay.

Acadex assists the controller, on the basis of the information available, with assessing and documenting the incident and with the resulting statutory obligations.

Security incidents are handled through a documented procedure assigning responsibilities for technical containment, investigation, documentation, assessment and communication.

11. Testing, assessment and evaluation of effectiveness

The effectiveness of safeguards is assessed through automated technical checks and operational controls.

These include in particular:

- automated unit tests
- selected User Acceptance Tests
- code review in the pull-request process
- automated checks of deletion logic within the database schema
- review of new Product Metrics for data minimisation and exclusion of content data
- technical error analysis using limited log data
- regular rotation of secrets and credentials
- multi-factor authentication for administrative platform access

- review of backup recoverability; a regular PostgreSQL restore-testing procedure is to be formalised as described in section 7.4

The technical and organisational measures are also reviewed on an event-driven basis following material changes to the system architecture or relevant security incidents. A fixed recurring formal review process is still being established.

Regular review interval: [TO BE COMPLETED]

12. Document maintenance

These Technical and Organisational Measures are versioned separately from the main agreement because technical safeguards may evolve during operation.

Changes are documented.

The version applicable to a particular contractual relationship is determined by the Data Processing Agreement and the incorporation of this annex agreed there.

Note: This working draft does not replace individual legal advice. The final measures and wording should undergo legal review before incorporation into customer agreements.