

Data Privacy Policy

Detailed policy for client data, AI inputs and outputs, embeddings, providers, retention, deletion and data requests

Core privacy commitment

KOMGIS will not treat client data as unrestricted raw material for AI experimentation. Client documents, prompts, retrieved context, AI outputs, embeddings, vector metadata and workflow records are governed data. Processing must be limited to an approved purpose, the minimum data needed, approved providers and storage locations, defined retention, and traceable deletion.

Policy owner: Privacy / AI Governance / Technical Leadership

Version: 1.0

Effective date: 8 July 2026

Classification: Internal Policy

Document Control

Field	Value
Document title	KOMGIS Data Privacy Policy
Document owner	Privacy / AI Governance / Technical Leadership
Version	1.0
Effective date	8 July 2026
Review frequency	At least annually and after defined trigger events
Approved by	[Executive approver / role]
Classification	Internal Policy
Applies to	KOMGIS employees, contractors, AI and software systems, client implementations and approved third parties
Related policies	Responsible AI and AI Management Policy; AI Security Policy; AI Incident Response Plan

Status and legal note

This policy is an internal operational framework for privacy and client-data handling in KOMGIS software and AI systems. It is informed by Regulation (EU) 2016/679 (GDPR), the Serbian Law on Personal Data Protection (Official Gazette of the Republic of Serbia No. 87/2018), relevant European Data Protection Board guidance, and KOMGIS's AI governance framework. It is not a public privacy notice, a data processing agreement, legal advice, or a determination that one rule applies identically in every jurisdiction. The responsible project team must determine KOMGIS's role, the applicable law, contractual instructions and system-specific requirements for each processing activity.

Policy Statement

KOMGIS designs AI and software systems so that privacy controls follow the data through ingestion, storage, retrieval, model processing, output generation, logging, backup and deletion. Privacy obligations do not stop when a document is converted into chunks, an embedding is created, context is sent to a model provider, or an AI output is generated.

KOMGIS will process personal data and confidential client data only for an approved purpose and under an identified legal, contractual or client-instruction basis. KOMGIS will minimize the data used, limit access, isolate clients, review providers, document processing locations, define retention, support data requests and propagate deletion to derived AI data.

The minimum internal rule is that prompts, model outputs, embeddings, vector metadata and AI-generated summaries are not presumed anonymous, harmless or outside privacy governance merely because they are technical or derived artifacts.

Contents

1. Purpose
2. Scope
3. Governance Basis and Policy Hierarchy
4. Definitions
5. Privacy Roles, Controller and Processor Responsibilities
6. Privacy Principles
7. Data Inventory and Data Processing Map
8. Data Classification and Sensitive Data
9. Lawful Basis, Client Instructions and Purpose Limitation
10. Client Data Collection and Ingestion
11. AI Inputs, Prompts and Retrieved Context
12. AI Outputs, Inferences and Derived Data
13. Embeddings and Vector Databases
14. RAG and Permission-Aware Data Processing
15. Model Providers, Vendors and Subprocessors
16. Model Training, Fine-Tuning and Evaluation Data
17. Data Minimization, Redaction and De-identification
18. Development, Testing, Demonstration and Support Data
19. Access Control, Confidentiality and Client Isolation
20. Security of Processing and Encryption
21. Logging, Monitoring, Telemetry and Analytics
22. International and Cross-Region Data Transfers
23. Retention and Storage Limitation
24. Deletion, Erasure and Data-Lineage Propagation
25. Backups, Recovery and Delayed Deletion
26. Client and Data Subject Requests
27. Automated Decision-Making, Profiling and Human Review
28. Transparency and Privacy Notices
29. Privacy Risk Assessment and DPIA
30. Personal Data Breach and AI Incident Coordination
31. Records, Audit and Accountability
32. Personnel, Training and Confidentiality
33. Exceptions, Nonconformity, Review and Enforcement
Appendix A. AI Data Lifecycle Processing Map
Appendix B. Data Classification and Handling Matrix
Appendix C. Baseline Retention Schedule
Appendix D. Deletion Propagation Checklist
Appendix E. Data Request Workflow
Appendix F. Provider Privacy Review Checklist
Appendix G. Required Privacy Records
Appendix H. Reference Documents

1. Purpose

This policy defines how KOMGIS protects personal data and confidential client data throughout the lifecycle of AI-enabled and conventional software systems. It establishes mandatory controls for data collection, ingestion, transformation, retrieval, model processing, output generation, storage, sharing, logging, retention, deletion and data requests.

The policy is intended to ensure that privacy is built into KOMGIS systems and operating procedures rather than handled only through contractual language or end-of-project review.

- identify what data is processed, where it moves, who can access it and why;
- distinguish KOMGIS responsibilities when acting as a controller, processor or other legally defined role;
- limit AI inputs and retrieved context to data necessary for the approved task;
- govern prompts, outputs, embeddings and vector records as data assets with defined owners and lifecycles;
- prevent cross-client exposure and unauthorized reuse of client data;
- approve model providers and subprocessors before confidential or personal data is sent to them;
- define retention and deletion rules for both source and derived AI data;
- support access, correction, restriction, export, objection and deletion requests according to applicable law and contract;
- maintain evidence that privacy controls are designed, implemented, tested and reviewed.

Mandatory privacy boundary

A model, prompt, embedding or vector index is not a privacy boundary. KOMGIS must enforce privacy through application logic, access control, tenant isolation, provider configuration, data minimization, encryption, retention controls and deterministic deletion workflows.

2. Scope

This policy applies to personal data, confidential client data and derived AI data processed by or on behalf of KOMGIS in connection with internal operations, software development, managed services, hosted products and client-specific implementations.

It applies to the following systems and processing environments, whether operated by KOMGIS or an approved third party:

- AI assistants, chatbots, voice interfaces and AI-enabled avatars;
- retrieval-augmented generation (RAG), enterprise search and document intelligence platforms;
- document ingestion, OCR, parsing, chunking, summarization, extraction and classification pipelines;
- AI agents, workflow automation and tool-using systems;
- model gateways, APIs, hosted models, local models and model-provider integrations;
- embedding services, vector databases, search indexes and caches;
- application databases, object stores, file systems, message queues and integration stores;
- development, test, staging, demonstration, production, support and incident environments;
- logs, analytics, monitoring, evaluation datasets and backups.

The policy applies to employees, contractors, developers, administrators, support staff, project managers, AI system owners and approved third parties with access to KOMGIS systems or client data.

This policy does not replace a public privacy notice or client-specific DPA. Those documents must accurately reflect the actual processing implemented under this policy.

3. Governance Basis and Policy Hierarchy

KOMGIS's privacy governance is based on the principles of lawfulness, fairness, transparency, purpose limitation, data minimization, accuracy, storage limitation, integrity, confidentiality and accountability. Where GDPR applies, these

principles are interpreted consistently with Regulation (EU) 2016/679. Where Serbian data protection law applies, the responsible owner shall assess the corresponding obligations under the Serbian Law on Personal Data Protection.

For AI systems, this policy must be read together with KOMGIS's Responsible AI and AI Management Policy. That policy already treats prompts, outputs, embeddings, vector metadata and workflow records as client data when generated from or connected to client content. This policy defines the more detailed privacy controls for those data categories.

Document	Primary role
Data Privacy Policy	Privacy roles, personal and client data, AI inputs/outputs, embeddings, providers, retention, deletion and requests.
Responsible AI and AI Management Policy	AI governance, lifecycle, risk, oversight, providers, RAG, agents, testing and incidents.
AI Security Policy	Technical protection against unauthorized access, prompt injection, data leakage, unsafe tool use and insecure development.
AI Incident Response Plan	Containment, evidence preservation, impact assessment, communication, corrective action and reactivation.
Client contract / DPA	Client-specific processing instructions, retention, residency, subprocessor and return/deletion commitments.
System records	Data Processing Map, AI System Inventory, retention schedule, provider review, DPIA and request records.

Where a client contract, applicable law or approved system requirement sets a stricter privacy control, the stricter control applies. Conflicts must be escalated before processing continues.

4. Definitions

Term	Policy meaning
Personal data	Information relating to an identified or identifiable natural person.
Client data	Information provided by, generated for, or connected to a client, including source files, records, prompts, retrieved context, outputs, embeddings, vector metadata and workflow records.
Derived AI data	Data created from source or client data through parsing, chunking, embedding, extraction, summarization, inference, classification or other AI processing.
Controller	The party that determines the purposes and means of personal data processing, as assessed under applicable law.
Processor	A party that processes personal data on behalf of a controller and under documented instructions, as assessed under applicable law.
Subprocessor	A processor engaged by another processor to perform processing on behalf of the controller.
Data subject	The natural person to whom personal data relates.
Special category data	Personal data subject to enhanced legal protection, such as certain health, biometric, genetic, political, religious, trade union or sex-life data where applicable law defines those categories.
Prompt / AI input	User, application, system or workflow content supplied to an AI model, including retrieved context.
AI output	Text, audio, image, structured data, recommendation, classification, extraction, tool request or other content generated by an AI system.
Embedding	A numerical representation of content used for similarity, retrieval or related machine-learning operations.
Vector record	An embedding and its associated identifiers, metadata, access-control attributes and source references.
RAG	A system that retrieves information from authorized data sources and provides selected context to an AI model.

Term	Policy meaning
Pseudonymization	Processing that separates data from direct identifying information so the data cannot be attributed to a person without additional information kept separately and protected.
Anonymization	A process intended to make a person no longer identifiable by means reasonably likely to be used, assessed in context and not assumed merely because direct identifiers were removed.
Deletion propagation	A controlled process that applies deletion or disabling to source data and linked derived data, indexes, caches and other active copies according to the data lineage.
Legal hold	A documented requirement to preserve specified data because of a legal, regulatory, investigation or dispute obligation.

5. Privacy Roles, Controller and Processor Responsibilities

KOMGIS shall determine and document its privacy role for each material processing activity. The same organization may be a controller for one activity and a processor for another. The use of an AI model provider does not by itself determine KOMGIS's legal role.

5.1 When KOMGIS acts as a processor

Where KOMGIS processes personal data on behalf of a client, KOMGIS shall process the data only under documented client instructions and the applicable contract or DPA, except where law requires otherwise. The project team shall maintain a processing map that is consistent with those instructions.

- ensure personnel with data access are subject to appropriate confidentiality obligations;
- use only approved subprocessors and follow contractual notification or approval requirements;
- implement appropriate technical and organizational measures;
- assist the client with data subject requests, DPIAs, security obligations and breach assessment as required by contract and law;
- make relevant processing and control information available for approved audit or assurance activities;
- return or delete personal data at the end of services according to the contract, subject to documented legal retention requirements;
- escalate instructions that appear unlawful, outside contract scope or materially inconsistent with the approved system design.

5.2 When KOMGIS acts as a controller

Where KOMGIS determines the purpose and means of processing, KOMGIS shall identify the applicable lawful basis, provide required transparency, maintain records of processing, support data subject rights and assess whether a DPIA or other privacy review is required.

5.3 Role ambiguity

If the role is unclear, the project shall not rely on informal labels such as "data owner" or "service operator". Privacy / Legal shall assess the actual decisions made about purpose, essential processing means, data use and onward disclosure. The result shall be recorded before high-impact or large-scale personal data processing begins.

6. Privacy Principles

Principle	KOMGIS operational rule
Lawfulness, fairness and transparency	Processing must have an identified basis and must not use hidden, deceptive or materially unexpected data practices.
Purpose limitation	Data is used for the approved purpose and compatible documented uses. New AI capabilities do not automatically authorize new purposes.
Data minimization	Only data necessary for the task is collected, retrieved, sent to a model or retained.

Principle	KOMGIS operational rule
Accuracy	Material personal data and AI-derived records must be correctable. Known errors must not be silently propagated into downstream workflows.
Storage limitation	Source and derived data receive defined retention periods and deletion triggers.
Integrity and confidentiality	Access control, encryption, client isolation, secure development and incident response protect data.
Accountability	KOMGIS maintains evidence of decisions, controls, assessments, provider reviews, requests and deletion.
Privacy by design and default	The default configuration minimizes exposure and does not require users to discover privacy protections themselves.

AI-specific interpretation

The fact that a model technically accepts a large context window or a provider permits long retention does not mean KOMGIS may use those capabilities by default. The approved purpose and minimum necessary processing determine the permitted data flow.

7. Data Inventory and Data Processing Map

Every material system that processes client or personal data shall maintain a Data Processing Map. The map is the operational source of truth for how data moves through the system and must be updated when architecture, provider, region, retention or data categories materially change.

The map shall identify, at minimum:

- system, client or service name and accountable owner;
- KOMGIS role and client/controller relationship where relevant;
- processing purpose and intended AI or software function;
- categories of personal data, confidential data and data subjects;
- source systems and ingestion methods;
- parsing, OCR, extraction, chunking and transformation steps;
- where embeddings, vector metadata, indexes and caches are created;
- model providers, model gateway, subprocessors and infrastructure providers;
- data sent to each provider, including prompt, retrieved context and output handling;
- storage systems, processing regions and transfer mechanisms where applicable;
- users, roles, services and support functions with access;
- retention period, deletion trigger, deletion propagation path and backup expiry;
- logging, telemetry and analytics content;
- client-specific restrictions, residency or no-training requirements.

The Data Processing Map must reflect the implemented architecture. A sales diagram or generic cloud drawing is not sufficient where it omits embeddings, temporary extraction files, provider-side processing, caches, logs or backups.

7.1 Data lineage identifiers

Systems that create derived AI data should maintain stable identifiers that link a source document or record to extracted text, chunks, embeddings, vector records and generated artifacts where practical. These identifiers support source traceability, permission changes, request handling and deletion propagation.

8. Data Classification and Sensitive Data

Data shall be classified according to the highest sensitivity reasonably present in the dataset or processing context. The classification controls which providers, environments, users and processing patterns are permitted.

Class	Examples / meaning	Minimum handling
Public	Intended for public release; still subject to integrity and copyright controls.	Approved public AI tools may be used only where separately authorized.
Internal	KOMGIS business information not intended for public release.	Approved enterprise tools; no personal-account AI use.
Confidential client	Client documents, records, business data, contracts, source code and workflow content.	Approved client-scoped systems and providers; tenant isolation and restricted support access.
Personal data	Information relating to an identified or identifiable natural person.	Document role, purpose, lawful basis/instructions, retention and rights handling.
Sensitive / regulated	Special category, criminal-offence, highly confidential or domain-regulated data.	Enhanced review, strict provider and region controls, minimum necessary data and formal risk assessment.
Secrets / credentials	Passwords, API keys, private keys, access tokens and encryption keys.	Must not be placed in prompts or normal model context; managed through approved secret systems.

8.1 Special category and highly sensitive data

Processing special category, biometric, genetic, health, criminal-offence or similarly protected data requires a documented assessment of the applicable legal condition, necessity, provider suitability, access restrictions and retention. The project shall not assume that a general contract or client request is sufficient without confirming the required legal and contractual basis.

Sensitive data shall not be added to general-purpose evaluation datasets, demonstration environments, prompt libraries or reusable fine-tuning corpora without specific written approval and a separate privacy assessment.

9. Lawful Basis, Client Instructions and Purpose Limitation

Before KOMGIS processes personal data as a controller, the processing purpose and applicable lawful basis shall be documented. Potential bases may include performance of a contract, compliance with a legal obligation, legitimate interests or consent, depending on the actual context. Consent shall not be used as a default label where it is not freely given, specific, informed and capable of being withdrawn as required by law.

When KOMGIS acts as a processor, the client's documented instructions and the DPA define the permitted processing purpose. KOMGIS shall not independently reuse the client's personal data for unrelated product training, marketing, profiling or model improvement merely because the data is technically available.

9.1 New purpose review

A new AI capability, model or workflow may change the purpose or nature of processing. Before reusing existing data for a materially new purpose, the system owner shall assess:

- whether the new use is within existing client instructions and privacy notices;
- whether the data subjects would reasonably expect the use;
- the relationship between the original and proposed purpose;
- the sensitivity of the data and potential consequences to individuals;
- whether a new lawful basis, notice, consent, contract amendment or DPIA is required;
- whether data minimization or de-identification can achieve the objective with less risk.

No function-creep rule

A system that was approved to answer questions from client documents is not automatically approved to build employee profiles, score individuals, create training datasets or send data to a new model provider. The new use must be assessed and approved.

10. Client Data Collection and Ingestion

Client data shall be collected or ingested only from authorized sources and for the approved service purpose. The ingestion design shall make clear who can add data, what sources are permitted and how unauthorized or accidental ingestion is handled.

10.1 Approved ingestion sources

- client uploads through authenticated application interfaces;
- approved API, database, file-store, email, CRM, ERP or document-management integrations;
- client-provided migration datasets or controlled import jobs;
- public sources only where the use is lawful, documented and within the intended system purpose.

10.2 Ingestion controls

- associate each item with the correct client, workspace, project or tenant before processing;
- capture a source identifier and, where practical, ingestion timestamp and source system;
- limit accepted file types, size and processing behavior according to the security policy;
- treat hidden text, comments, annotations, OCR layers, spreadsheet hidden sheets and document metadata as potentially sensitive data;
- avoid broad "sync everything" integrations when a limited folder, field set or record subset is sufficient;
- document whether deleted or permission-revoked source items are automatically removed from downstream AI stores;
- detect and resolve duplicate ingestion where duplication would create unnecessary retention or inconsistent deletion.

KOMGIS shall not request clients to upload more personal data than is necessary for implementation or troubleshooting. Support channels shall not be used as uncontrolled repositories for production exports or document collections.

11. AI Inputs, Prompts and Retrieved Context

Prompts and model inputs may contain personal data, confidential business information and retrieved document content. They shall be treated as governed processing records, not disposable text.

11.1 Input minimization

Before sending data to a model, the application should select the minimum fields, excerpts, records or chunks necessary for the task. Entire document libraries, full database rows with unused columns, or complete conversation histories shall not be sent merely for implementation convenience.

- retrieve authorized context first, then construct the prompt from the selected context;
- remove irrelevant identifiers and sensitive fields where they do not contribute to the task;
- use structured inputs with explicit fields where this reduces accidental data exposure;
- limit conversation memory to the scope and duration needed for the user experience;
- do not include secrets, credentials or private keys in model context;
- do not use production prompts containing client data as public examples, screenshots or support samples.

11.2 System prompts and client data

System prompts should contain policy, task and behavior instructions, not client secrets or large amounts of personal data. Persistent system instructions are subject to code and configuration controls and should not be used as a substitute for secure data storage.

11.3 Conversation context and memory

Conversation history retained for continuity shall have a defined scope, retention and deletion behavior. Cross-user or cross-client memory is prohibited unless the feature is specifically designed, authorized and permission-controlled. A user must not receive another user's conversation content through model memory, caching or shared thread state.

12. AI Outputs, Inferences and Derived Data

AI outputs are governed according to their content, source and business use. An output may contain copied source material, inferred personal data, sensitive conclusions, fabricated personal data or new structured records. The fact that the content was generated by AI does not remove privacy obligations.

12.1 Output handling rules

- classify and store outputs according to the data they contain and the workflow they support;
- do not expose output to users who could not access the underlying client or case context;
- avoid creating persistent output records when a transient answer is sufficient;
- validate structured outputs before writing them to client systems;
- mark or trace AI-generated content where the workflow requires human review or source verification;
- allow material personal data errors to be corrected and prevent known incorrect output from being repeatedly reused;
- review outputs that infer health, employment, financial, insurance, legal, identity or similarly sensitive attributes before use in a material workflow.

12.2 Hallucinated or inaccurate personal data

If an AI system generates false personal data or attributes it to a real person, KOMGIS shall not treat the content as a reliable record. Where the output has been stored, communicated or used in a material process, the responsible owner shall correct or remove it, assess downstream impact and consider whether incident or data-request processes are triggered.

12.3 Output reuse

Client-specific outputs shall not be added to general prompt libraries, training datasets, demonstrations or shared evaluation corpora without explicit approval and documented privacy review. Output reuse must respect the original data source, client instructions and data subject rights.

13. Embeddings and Vector Databases

KOMGIS treats embeddings and vector records as derived client data when created from client content. KOMGIS will not presume that an embedding is anonymous, non-sensitive or safe for unrestricted sharing merely because the original text is not directly readable from the vector value.

13.1 Required vector record metadata

Production vector records should contain sufficient metadata to enforce access and lifecycle controls. Depending on architecture, this may include:

- client or tenant identifier;
- workspace, project, field, section or equivalent authorization scope;
- source document or record identifier;
- chunk or segment identifier;

- permission or access-control attributes required for filtering;
- source version or ingestion batch where relevant;
- retention or lifecycle state, including deleted, disabled or tombstoned status.

13.2 Vector-store privacy controls

- use tenant-aware namespaces, collections, partitions or metadata filters consistent with the application authorization model;
- apply client and user permission filtering before retrieved content enters the model prompt;
- restrict administrative and direct database access;
- protect vector stores and metadata in transit and at rest according to the security policy;
- do not combine client embeddings into a shared research or training corpus without explicit approval;
- test cross-client isolation and deletion behavior before production;
- rebuild, delete or disable embeddings when the source is deleted, access is revoked or the source is materially replaced.

Embedding lifecycle rule

The lifecycle of a vector record normally follows its source. Deleting a source document while leaving its active chunks and embeddings retrievable is a privacy and data-governance failure unless a documented retention exception applies.

14. RAG and Permission-Aware Data Processing

RAG systems create a privacy risk when retrieval is broader than the user's authorization. KOMGIS shall enforce source authorization before model context is constructed. The model must not be asked to hide or ignore content that the retriever should never have returned.

14.1 Minimum RAG privacy controls

- authenticate the user or calling service;
- resolve the active client, tenant, workspace and user authorization context;
- filter the candidate data set before or during retrieval using server-side permission rules;
- apply object-level or document-level authorization where the underlying system requires it;
- include only authorized chunks in prompt context;
- maintain source and chunk identifiers for citations, investigation and deletion;
- propagate source permission changes to retrieval indexes and caches;
- test unauthorized document queries, guessed identifiers and cross-client semantic similarity scenarios;
- treat retrieved content as untrusted data and do not allow document instructions to change privacy rules or export permissions.

14.2 Permission change propagation

When a user loses access to a project, section, case or document, the RAG layer shall stop returning the affected content without waiting for model retraining. Permission updates shall be enforced through deterministic authorization and retrieval filtering. Where search indexes cache permission metadata, update latency and failure handling must be documented and tested.

14.3 Citations and privacy

Source citations improve traceability but may themselves reveal file names, case identifiers or document existence. Citation interfaces shall display only source information the user is authorized to see. A citation link must not become an alternate path around application authorization.

15. Model Providers, Vendors and Subprocessors

Personal or confidential client data may be sent only to approved model providers, cloud services and subprocessors for the approved use case. Personal AI accounts, unknown public endpoints and unapproved model proxies shall not be used for client confidential data.

15.1 Provider privacy review

Before approval, KOMGIS shall assess the provider and specific service configuration. Review shall consider:

- exact product, API or model service being used;
- categories of data sent and whether special category or regulated data may be processed;
- provider role and contractual data-processing terms;
- whether prompts, outputs or files are used for model training, service improvement or human review;
- default and configurable request retention, abuse-monitoring retention and deletion options;
- processing and storage regions, data residency controls and cross-border transfer mechanisms;
- subprocessor list and change-notification process;
- security documentation, access controls, encryption and incident commitments;
- availability of enterprise controls such as zero-data-retention or reduced-retention modes where relevant;
- client restrictions, sector requirements and contractual approval obligations;
- ability to support data requests, deletion and incident investigation;
- provider exit, model change and continuity risk.

15.2 Configuration must match approval

Approval of a provider name does not approve every product, region or account type offered by that provider. The deployed account, endpoint, retention setting and training-use configuration shall match the approved provider record. Consumer or personal accounts shall not be substituted for approved enterprise services.

15.3 Provider changes

A new provider, model gateway, hosting region or materially different provider retention configuration is a controlled privacy change. The Data Processing Map, Vendor/Subprocessor List, transfer assessment, client requirements, retention schedule and evaluation evidence shall be reviewed before production use.

16. Model Training, Fine-Tuning and Evaluation Data

KOMGIS's default position is that client data is not used to train or fine-tune KOMGIS or third-party models unless the client has explicitly approved the use and the purpose, data scope, legal basis or processing instructions, retention, security and deletion limitations are documented.

16.1 Training and fine-tuning approval requirements

- document the training purpose and why existing models, prompting or RAG are insufficient;
- identify exact datasets, source clients and personal data categories;
- record client authorization and applicable legal basis or processing instruction;
- apply minimization, redaction, pseudonymization or anonymization where feasible;
- separate the approved dataset from ordinary production data and maintain data lineage;
- evaluate whether data subject rights and deletion can be honored after model training;
- document model-retention, checkpoint and derivative-data implications;
- approve the provider and region used for training;
- define who may access the dataset and how it will be deleted or archived.

16.2 Evaluation and red-team datasets

Evaluation datasets are subject to the same privacy principles as production data. KOMGIS should use synthetic, public, anonymized or specifically approved examples whenever possible. Production conversations, client documents and incident samples shall not be copied into reusable evaluation corpora without documented approval and retention.

16.3 Provider service improvement

Where a model or AI provider may use customer content for service improvement, model training, abuse review or human inspection, that behavior must be identified during provider review. KOMGIS shall use available enterprise opt-out, no-training or reduced-retention controls where required by the approved use case and client commitments.

17. Data Minimization, Redaction and De-identification

KOMGIS shall reduce data exposure before, during and after AI processing. Minimization is preferred over relying solely on contractual restrictions after unnecessarily broad data has already been transferred.

17.1 Minimization techniques

- select only required documents, records, columns and fields;
- retrieve only the top authorized chunks needed for the task;
- truncate long conversation history when old context is no longer needed;
- replace direct identifiers with internal identifiers where the task does not require names;
- redact secrets and unnecessary sensitive attributes before provider processing;
- use local preprocessing or deterministic extraction to avoid sending full documents when only a limited field set is required;
- store metadata rather than full prompt or output content in logs where possible.

17.2 Pseudonymization

Pseudonymized data remains governed data. Re-identification information shall be stored separately and protected with access controls. Pseudonymization may reduce risk but does not automatically remove the processing from data-protection obligations.

17.3 Anonymization

KOMGIS shall not label a dataset or AI artifact anonymous solely because names, email addresses or obvious identifiers were removed. The assessment shall consider singling out, linkability, inference, external data and the technical and organizational means reasonably available in context. Where material uncertainty remains, the data shall continue to be treated as personal or confidential data.

De-identification rule

Removing a person's name from a prompt, embedding source or evaluation sample is not enough when the remaining facts, identifiers or context can still identify the person. The practical re-identification risk must be assessed.

18. Development, Testing, Demonstration and Support Data

Development and testing shall use synthetic, dummy, public or appropriately de-identified data by default. Real production client data may be used only where necessary, approved and technically restricted.

18.1 Production data in non-production environments

Before production data is copied to development, test, staging or demonstration, the owner shall document why synthetic data is insufficient and approve the minimum data subset. The copy shall have a deletion date and shall not become a permanent test fixture.

- use separate accounts, projects and credentials for non-production environments;

- restrict developer access to the approved dataset and environment;
- remove or mask unnecessary personal and sensitive fields;
- prevent non-production data from being sent to unapproved model providers or analytics tools;
- delete temporary exports, local copies and test indexes after the approved task;
- do not use client data in public demos, screenshots, conference material or sales examples without specific approval.

18.2 Support and troubleshooting

Support personnel shall request the smallest evidence needed to diagnose an issue. Full database exports, document collections or complete prompt histories should not be requested when a request identifier, redacted example or metadata log is sufficient. Time-bound production access shall be used where practical and logged for sensitive systems.

18.3 Local devices

Production client data shall not be copied to unmanaged personal devices, personal cloud storage or consumer AI tools. Approved local development involving client data requires documented controls, encryption and deletion after use.

19. Access Control, Confidentiality and Client Isolation

Privacy depends on preventing access by persons and systems that do not need the data. KOMGIS shall apply least privilege and client isolation across application databases, file storage, vector stores, logs, analytics, provider accounts and support processes.

19.1 Access control requirements

- authenticated access unless data is explicitly intended to be public;
- role-based and object-level authorization enforced by backend services;
- tenant-aware database queries and file-access checks;
- permission-aware vector retrieval and source access;
- client-scoped service credentials and integrations where practical;
- least-privilege model-provider API keys and cloud roles;
- administrative access limited to named authorized personnel;
- periodic access review for high-sensitivity systems;
- prompt revocation and offboarding when employment, contractor or client access changes;
- logging of material administrative or production-data access.

19.2 Client isolation

One client's users, agents, service accounts, document retrievers, vector records, prompts, outputs and credentials shall not access another client's data. Shared infrastructure is permitted only where logical isolation is designed, implemented and tested at each relevant data layer.

Mandatory isolation rule

It is prohibited to retrieve broadly across clients and rely on the AI model to suppress unauthorized content. Authorization must determine the eligible data set before the model receives context.

20. Security of Processing and Encryption

KOMGIS shall implement technical and organizational measures appropriate to the sensitivity, scale, processing context and risk of the personal and client data processed. Privacy and security controls shall be coordinated with the AI Security Policy.

20.1 Minimum technical protection

- encrypt data in transit using approved secure protocols;
- encrypt confidential and personal data at rest where supported by the platform and architecture;
- manage encryption keys and secrets through approved systems and restrict access;
- protect storage buckets, databases, vector stores and provider accounts from public exposure;
- maintain secure software development, dependency and vulnerability controls;
- validate file ingestion and temporary-processing pipelines;
- separate sensitive production data from public test and demonstration systems;
- maintain backup, recovery and deletion procedures;
- monitor for unusual access, data export and cross-client indicators.

20.2 Secrets are not model context

Passwords, API keys, access tokens, private keys, encryption keys and similar secrets must not be placed in prompts, retrieved context, AI memory or training data. If a secret is exposed to a model or provider, the event must be reported and the secret rotated or revoked according to the incident process.

21. Logging, Monitoring, Telemetry and Analytics

Logs and observability systems can become secondary stores of client and personal data. KOMGIS shall design logging to collect the evidence needed for security, reliability, billing, debugging and traceability while minimizing content exposure.

21.1 Logging preference

Where possible, logs should contain metadata such as request identifier, timestamp, tenant, system version, model identifier, source identifiers, tool status, validation result and error class rather than full prompts, documents or outputs.

Log type	Typical content	Privacy rule
Audit metadata	User/service ID, tenant, timestamp, system/version, action, source IDs, approval status.	Restricted; may have longer retention for security and accountability.
Prompt/output content logs	Full or partial model inputs and outputs.	Disabled or minimized by default; highly restricted; shortest practical retention.
Debug logs	Errors, stack context, processing diagnostics.	Redact content and secrets; short technical retention.
Security logs	Authentication, authorization failures, suspicious export/retrieval activity.	Restricted security access; retention based on threat and compliance needs.
Telemetry / analytics	Usage and performance metrics.	No client content or personal identifiers unless approved and documented.
Provider request records	Provider request IDs, model, region and status.	Retain only as needed for support, billing, incident investigation or audit.

21.2 Sensitive content in logs

Where full prompt or output content logging is necessary, the system owner shall document the purpose, access restrictions and retention. Support or engineering personnel shall not use content logs for curiosity, informal product review or unrelated analysis.

21.3 Third-party analytics

Client document text, prompts, outputs and direct identifiers shall not be sent to general telemetry or product analytics services unless the service is approved for that data and the processing is documented. Event names, URLs and error messages shall be reviewed for accidental data leakage.

22. International and Cross-Region Data Transfers

KOMGIS shall document where personal and client data is stored and processed, including model-provider inference regions where relevant. Client residency requirements must be captured before deployment and reflected in architecture and provider configuration.

22.1 Transfer assessment

Where personal data is transferred internationally or made accessible from another jurisdiction, KOMGIS shall assess the applicable transfer rules and contractual requirements. Depending on the governing law, this may require an adequacy basis, standard contractual clauses, another recognized transfer mechanism, supplementary measures or a documented transfer assessment.

22.2 Data residency controls

- record approved storage and processing regions in the Data Processing Map;
- configure cloud and model-provider regions according to the approved design;
- do not move a workload to another region for convenience without privacy and client review;
- consider support access and subprocessor locations, not only the primary database region;
- review region implications when changing model provider, vector database, backup service or analytics platform.

Where a provider cannot provide sufficient information about processing locations or transfer terms for the intended data, KOMGIS shall not send the data until the risk is resolved or an approved alternative is selected.

23. Retention and Storage Limitation

Client and personal data shall be retained only as long as needed for the approved purpose, client contract, legal requirement, security need or documented audit requirement. Retention applies to both source data and derived AI data.

23.1 System retention schedule

Every material AI or client-data system shall maintain a retention schedule identifying the data type, purpose, retention period or event-based trigger, deletion method, backup treatment and owner. "Keep indefinitely" is not an acceptable default unless a specific legal or contractual reason is documented.

23.2 Retention criteria

- duration of the client service or business workflow;
- client contractual retention and return/deletion commitments;
- legal or regulatory recordkeeping requirements;
- need to preserve security, fraud, audit or incident evidence;
- whether the same purpose can be achieved with metadata rather than content;
- risk created by keeping prompts, outputs, embeddings and chat history;
- provider-side retention that KOMGIS cannot technically control;
- backup rotation and disaster-recovery requirements.

23.3 Retention defaults

Appendix C defines KOMGIS baseline retention expectations where a system-specific schedule has not set a stricter or contractually different period. The owner may approve a different period only when the purpose and legal or contractual reason are documented.

Derived-data retention rule

Chunks, extracted text and embeddings normally inherit the lifecycle of the source. AI outputs and chat history require their own purpose-based retention because they may become business records or may be unnecessary transient content.

24. Deletion, Erasure and Data-Lineage Propagation

Deletion must address the real data lifecycle. Removing a source file from the application interface is not sufficient when extracted text, chunks, embeddings, cached context, generated summaries or active indexes remain available.

24.1 Deletion triggers

- client or authorized user deletion of a source item;
- approved client offboarding or service termination;
- expiry of the defined retention period;
- validated data subject erasure request where erasure applies;
- withdrawal of consent where consent is the applicable basis and no other basis requires retention;
- correction or replacement of source data that requires stale derived data to be removed;
- privacy or security incident requiring containment or removal;
- manual administrative deletion approved through the defined process.

24.2 Deletion propagation scope

The deletion workflow shall evaluate all linked data stores and processing artifacts, including:

- uploaded files and object-storage versions;
- application database records;
- extracted text, OCR output and parser artifacts;
- chunks, embeddings and vector metadata;
- search indexes and cached retrieval results;
- AI-generated summaries, reports and workflow artifacts where deletion applies;
- chat history and conversation memory;
- temporary files, queues and integration caches;
- prompt/output logs where deletion applies and content can be located;
- provider-side stored files or datasets where the provider service maintains persistent copies;
- backups according to the backup-retention process.

24.3 Active retrieval disablement

A validated deletion or permission revocation affecting RAG content should remove or disable the affected content from active retrieval as quickly as technically practical. Systems may use tombstones or lifecycle flags to block retrieval immediately while background hard-deletion completes. Tombstoned content shall not be returned to users or models.

24.4 Deletion verification

Deletion jobs shall produce a status or audit record sufficient to show the source identifier, deletion request, systems evaluated, completion or exception status and timestamp. Verification should use deterministic identifiers and store queries rather than asking an AI model to decide whether data is deleted.

24.5 Exceptions and legal holds

Where data cannot be deleted because of a legal hold, regulatory requirement, dispute or other documented obligation, the data shall be restricted from ordinary use where appropriate. The exception, authority, scope and review date shall be recorded. KOMGIS shall not keep all client data indefinitely merely because a future dispute is theoretically possible.

25. Backups, Recovery and Delayed Deletion

Backups shall be protected against unauthorized access and retained according to the approved backup schedule. Because individual record deletion from immutable backups may be technically impractical, KOMGIS shall manage backup deletion through rotation, expiry and controlled restore procedures.

25.1 Backup deletion treatment

- document the backup retention period for each material data store;
- prevent expired backups from being retained indefinitely;
- limit access to backup systems and encryption keys;
- do not restore a deleted record into active service without reapplying deletion state;
- maintain a deletion ledger, tombstone list or equivalent mechanism where needed so that restore procedures re-execute deletions;
- ensure backup copies are not used as analytics, training or test datasets;
- document any contractually committed maximum backup-expiry period.

25.2 Restore procedure

After a restore, the technical owner shall verify that current tenant permissions, deleted-source states, legal holds and retention rules are reapplied before the restored system is returned to normal use. RAG indexes and embeddings rebuilt from a backup shall not reintroduce content that was deleted after the backup was created.

26. Client and Data Subject Requests

KOMGIS shall maintain a documented process to receive, verify, track and respond to client data requests and data subject rights requests. The process shall distinguish requests KOMGIS handles as controller from requests KOMGIS supports on behalf of a client controller.

26.1 Request types

- access to personal data or client data;
- correction or rectification;
- deletion or erasure;
- restriction of processing;
- data portability or structured export where applicable;
- objection to processing;
- withdrawal of consent where applicable;
- information about model providers, subprocessors, processing regions and retention;
- questions about automated decision-making, profiling or human review;
- client offboarding, return and deletion of client data.

26.2 Identity and authority verification

Before disclosing, correcting or deleting data, KOMGIS shall verify the identity and authority of the requester to a level proportionate to the request and risk. KOMGIS shall avoid collecting excessive new identity data merely to verify a request. Where a client administrator submits a request, the system or support team shall confirm the administrator's authority for the affected workspace or data.

26.3 Processor handling

Where KOMGIS acts as processor and receives a data subject request directly, KOMGIS shall promptly route the request to the relevant client/controller and shall not independently determine the response unless authorized or required by law. KOMGIS shall assist with locating, exporting, correcting, restricting or deleting data according to the DPA and applicable instructions.

26.4 Controller handling and deadlines

Where KOMGIS acts as controller, Privacy / Legal shall determine the applicable right, exceptions and response deadline. Where GDPR applies, the response process shall be designed around the one-month period in Article 12, subject to the lawful extension mechanism for complex or numerous requests. Requests shall be recorded on receipt and escalated early when scope is unclear.

26.5 Locating AI-related data

Request searches shall consider source records and derived AI data. Depending on the system, the search may need to evaluate prompts, chat history, outputs, extracted text, embeddings, vector metadata, model-provider stored files, audit records and business workflow records. The Data Processing Map and source identifiers shall guide the search.

26.6 Response quality

Responses must be accurate and understandable. KOMGIS shall not state that data was fully deleted when only the visible source was removed or when known active derived copies remain. Where deletion is delayed by backups, legal holds or provider limitations, the responsible owner shall describe the status consistently with applicable law and contract.

27. Automated Decision-Making, Profiling and Human Review

AI features that evaluate, score, rank, recommend or infer information about people may create privacy and fundamental-rights implications even when described as decision support. The system owner shall document whether the output materially influences legal, financial, employment, insurance, health, education, eligibility, safety or similarly important outcomes.

27.1 Solely automated decisions

KOMGIS shall not deploy solely automated decision-making that produces legal effects or similarly significant effects on a person without a specific assessment of the applicable legal conditions, safeguards, transparency and human intervention rights. The presence of a human in the workflow is not sufficient if the human routinely approves AI output without meaningful authority or evidence.

27.2 Profiling and inferred data

Systems that infer traits, risk levels, preferences, behavior or other personal attributes shall identify the data sources, purpose, accuracy risks and consequences of error. Sensitive inferences shall not be treated as facts merely because the model expresses them confidently.

27.3 Human review

Where meaningful human review is required, the reviewer shall have access to relevant source evidence, understand the system's limitations, be able to reject or override the output and have sufficient time and authority to do so. Review decisions should be recorded for high-impact workflows.

28. Transparency and Privacy Notices

KOMGIS shall provide or support accurate privacy information about data processing where required by law or contract. Public privacy notices, client notices and in-product disclosures must reflect the implemented processing rather than generic claims about AI.

28.1 Information to assess for disclosure

- identity and contact details of the relevant controller;
- processing purposes and applicable lawful basis where KOMGIS is controller;
- categories and sources of personal data;
- AI-assisted nature of the feature where relevant;
- categories of recipients, model providers or subprocessors as required;
- international transfer information where required;
- retention period or criteria;
- data subject rights and request channels;
- meaningful information about automated decision-making where applicable;
- role of human review in sensitive workflows.

28.2 Security transparency boundary

Transparency does not require KOMGIS to publish exploitable secrets, credentials, internal attack-detection rules or detailed security architecture. Privacy communications should be specific about processing while protecting information that would materially weaken security.

28.3 Model-provider transparency

When clients request information about model providers and subprocessors, KOMGIS shall use the approved Vendor/Subprocessor List and Model Provider Disclosure. Staff shall not guess based on the model name shown in code or a user interface.

29. Privacy Risk Assessment and DPIA

Privacy risk assessment shall be integrated into system design and AI risk management. A Data Protection Impact Assessment (DPIA) or equivalent formal assessment shall be considered before processing that is likely to result in high risk to individuals, particularly where new technologies, large-scale sensitive data, systematic monitoring, profiling or important decisions are involved.

29.1 Screening questions

- Does the system process special category, biometric, genetic, health or criminal-offence data?
- Does it evaluate, score, profile or infer information about people?
- Could an error materially affect rights, access, money, employment, insurance, health, safety or essential services?
- Is processing large scale or does it combine data from multiple sources?
- Does the system systematically monitor individuals or behavior?
- Is the processing materially unexpected or difficult for individuals to avoid?
- Are vulnerable persons, minors or employees materially affected?
- Does the system use a new provider, cross-border transfer, high-autonomy agent or novel AI technique?
- Can the system honor access, correction and deletion across source and derived data?
- Would a data breach, wrong-client retrieval or sensitive inference cause high impact?

29.2 DPIA content

A DPIA or formal privacy impact assessment should describe the processing, purpose, necessity and proportionality; assess risks to individuals; identify controls; evaluate residual risk; define owners and evidence; and record approval. Where required by applicable law, unresolved high residual risk shall be escalated for prior consultation or other legal review before processing begins.

29.3 AI model and anonymization assessment

Where a project relies on a claim that an AI model, dataset or derived artifact is anonymous, the assessment shall be documented and reviewed in context. KOMGIS shall take a conservative approach where model memorization, extraction, linkage or inference risks are not sufficiently understood.

30. Personal Data Breach and AI Incident Coordination

A privacy incident may also be an AI security or operational incident. Examples include wrong-client RAG retrieval, prompt leakage, model output containing another person's data, unauthorized provider processing, exposed embeddings, public storage, excessive logging or an agent sending personal data to an unauthorized recipient.

30.1 Immediate actions

1. contain the exposure and stop ongoing processing or retrieval where necessary;
2. preserve relevant logs, source IDs, prompts, provider request IDs, approval records and access evidence;
3. identify affected data, clients, individuals, systems, regions and recipients;

4. revoke credentials, disable tools, remove public access or isolate indexes as needed;
5. coordinate Security, Privacy, AI Governance, system owners and client owners;
6. assess whether the event is a personal data breach and whether notification duties apply;
7. support the client/controller with information required for its breach assessment where KOMGIS is processor;
8. correct the technical and process cause and verify controls before reactivation.

30.2 Regulatory and client notification

Where GDPR applies and KOMGIS is controller, the incident process shall support the supervisory-authority notification assessment required by Article 33, including the 72-hour notification period where notification is required. Where KOMGIS is processor, KOMGIS shall notify the relevant controller without undue delay after becoming aware of a personal data breach, consistent with applicable law and the DPA. Other jurisdictions or contracts may impose different or additional requirements.

30.3 Evidence and privacy

Incident preservation shall be targeted. KOMGIS shall not create unnecessary full copies of sensitive datasets merely for investigation. Where content evidence is required, access shall be restricted and the incident record shall define the retention and deletion of preserved evidence.

31. Records, Audit and Accountability

KOMGIS shall maintain privacy records proportionate to the processing and risk. Records must be accurate enough to demonstrate how the system actually handles data and to support audits, client reviews, incidents and data requests.

Required records may include:

- Data Processing Map and data-flow architecture;
- record of processing activities where applicable;
- client DPA and documented processing instructions;
- AI System Inventory and risk classification;
- data classification and retention schedule;
- Vendor/Subprocessor List and provider privacy reviews;
- international transfer assessment and region configuration;
- DPIA or privacy impact assessment;
- access review and production-data access records;
- data request register and completion evidence;
- deletion propagation and offboarding records;
- privacy incidents, breaches and corrective actions;
- training and confidentiality records;
- approved exceptions and legal holds.

31.1 Audit review

Internal audit or governance review should sample systems with client data, RAG, embeddings, external model providers or sensitive workflows. Review should compare the approved records against the implemented storage, provider calls, retention jobs, access controls and deletion behavior. A complete policy document is not evidence that the technical system follows the policy.

32. Personnel, Training and Confidentiality

Personnel with access to personal or client data shall understand the sensitivity of AI-related data and the limits of approved tools. Training shall be appropriate to the role and system risk.

Audience	Minimum privacy competence
----------	----------------------------

Audience	Minimum privacy competence
All staff	Approved AI tools, confidentiality, data minimization, no secrets in prompts, incident reporting and safe data sharing.
Developers	Data mapping, tenant isolation, vector-store privacy, prompt/output logging, provider configuration, deletion propagation and secure test data.
Support	Minimum necessary troubleshooting data, identity/authority checks, restricted production access and request escalation.
System owners	Role assessment, lawful purpose/client instructions, provider approval, retention, DPIA screening and risk acceptance.
Privacy / Legal	Applicable law, controller/processor roles, rights, transfers, DPIA, breach assessment and client/privacy communications.
Human reviewers	Sensitive inferences, correction of personal data, automation bias, evidence review and escalation.

Confidentiality obligations continue after access or engagement ends according to contract and applicable law. Departing personnel shall lose access to client systems, provider accounts, vector stores, logs and production support channels through the offboarding process.

33. Exceptions, Nonconformity, Review and Enforcement

33.1 Exceptions

An exception to this policy must be documented, time-bound and approved by the appropriate system, privacy, security or leadership owner according to risk. The exception record shall identify the requirement, reason, data affected, compensating controls, residual risk, owner and expiry or review date.

Convenience, model quality or implementation speed alone is not sufficient justification for broad client data collection, indefinite retention, cross-client reuse or sending data to an unapproved provider.

33.2 Nonconformity and corrective action

Privacy nonconformities may be identified through testing, audit, incidents, data requests, client complaints, support review or monitoring. KOMGIS shall contain the issue, identify affected systems and data, correct the immediate problem, determine root or contributing causes, assign corrective actions and verify effectiveness.

33.3 Policy review triggers

- new material AI system, RAG platform or agent capability;
- new model provider, subprocessor or significant processing region;
- material change to provider training-use or retention terms;
- new special category or regulated data processing;
- high or critical privacy or AI incident;
- significant client contractual or residency requirement;
- new or changed legal, regulatory or data-protection guidance;
- audit findings, recurring requests or deletion failures indicating a policy gap.

33.4 Enforcement

Violations may result in removal of data or AI tool access, suspension of a feature or integration, corrective training, project escalation, disciplinary action consistent with employment arrangements, or termination of third-party access. Good-faith privacy, security or AI concerns shall be reported through approved channels and must not be suppressed because of delivery pressure.

Appendix A. AI Data Lifecycle Processing Map

The following lifecycle should be used as a baseline when documenting an KOMGIS AI or document-intelligence system. Not every system uses every stage, but omitted stages should be intentionally marked not applicable.

Stage	Processing example	Required mapping fields
1. Source	Client document, database, API, email, CRM/ERP or user input.	Source owner; data categories; data subjects; purpose; permission source.
2. Ingestion	Upload, sync, API pull, import or migration.	Tenant assignment; source ID; region; temporary files; duplicate handling.
3. Extraction	OCR, parsing, field extraction, normalization.	Output location; sensitive fields; temporary retention; provider use.
4. Chunking	Source divided into retrieval segments.	Source/chunk IDs; permission inheritance; source version.
5. Embedding	Text converted to vector representation.	Embedding provider; data sent; region; vector lifecycle.
6. Vector storage	Embedding and metadata stored for retrieval.	Tenant isolation; authorization metadata; retention; encryption.
7. Retrieval	Authorized search selects chunks or records.	User/tenant context; filters; access checks; retrieval logs.
8. Prompt construction	System instructions plus user input and retrieved context.	Minimum context; redaction; conversation memory; secrets excluded.
9. Model processing	Prompt sent to local or external model.	Provider; model; training use; retention; region; request ID.
10. Output	Answer, extraction, classification, recommendation or tool request.	Sensitivity; validation; human review; storage and correction.
11. Workflow use	Output displayed, stored or used in downstream action.	Recipients; write actions; approvals; business-record retention.
12. Logging	Audit, error, security and usage records.	Metadata vs content; redaction; access; retention.
13. Backup	Protected recovery copy.	Backup scope; encryption; rotation; deletion replay.
14. Deletion	Source and linked data removed or disabled.	Propagation path; provider delete; backup expiry; verification evidence.

Appendix B. Data Classification and Handling Matrix

Class	Examples	AI/provider rule	Access	Retention	Deletion
Public	Approved public content.	Approved system.	Business need.	Purpose-based.	Normal deletion process.
Internal	Internal docs, non-public business data.	Approved enterprise services.	KOMGIS roles only.	Defined by system.	Delete at expiry/offboarding.
Confidential client	Client docs, source code, workflows, contracts.	Client-approved KOMGIS systems/providers.	Tenant-scoped least privilege.	Client/service schedule.	Propagate to chunks, embeddings and outputs as applicable.
Personal data	Names, emails, identifiers, behavior, records.	Approved processing map and provider.	Role/object authorization.	Purpose/legal basis or instructions.	Rights-aware deletion and exceptions.
Sensitive / regulated	Health, biometric, genetic, criminal, highly regulated.	Specifically approved provider and region.	Enhanced restricted access.	Minimum necessary and shortest justified.	Formal verification; legal hold handling.
Secrets	API keys, tokens, passwords, private keys.	Secret manager only.	Named service/admin access.	Rotate and expire.	Revoke/rotate; never retain in prompts.

Appendix C. Baseline Retention Schedule

The table below is KOMGIS's baseline where a system-specific retention schedule, client contract or applicable law does not require a different period. Owners should prefer shorter periods when the purpose can be met with less retention.

Data type	Baseline retention	Deletion trigger	Possible exception
Uploaded source files	Duration of active service/workflow plus up to 30 days for orderly deletion unless client schedule differs.	Client deletion, offboarding or expiry.	Legal hold or specific contractual retention.
Extracted text / OCR / chunks	Same lifecycle as source; temporary parser artifacts normally 24 hours or less.	Source deletion/replacement.	Documented audit need.
Embeddings / vector records	Same lifecycle as source.	Source deletion, permission withdrawal or re-index.	Documented retention exception; content must be non-retrievable if access revoked.
Transient model context	No application retention beyond request unless required for the feature.	Request completion.	Provider technical retention documented separately.
Chat history	90 days baseline where continuity is needed; shorter for sensitive use; client-defined where applicable.	Expiry, user/client deletion, offboarding.	Business-record or legal requirement.
AI outputs / summaries	90 days baseline for transient assistance; business workflow records follow the underlying case/process schedule.	Expiry or workflow deletion.	Legal, audit or contractual record.
Full prompt/output content logs	30 days or less by default when enabled.	Automatic log expiry.	Active incident or approved debugging case.
Debug / error logs	14 days baseline.	Automatic expiry.	Active incident investigation.
Security / audit metadata	12 months baseline.	Automatic expiry.	Longer regulatory, contractual or incident requirement.
Temporary exports / support copies	24 hours to 7 days depending on task; explicit deletion date required.	Task completion.	Active support case with approved extension.
Evaluation datasets with client data	Project-defined shortest period; must have explicit approval.	Evaluation completion.	Approved benchmark maintenance with documented basis.
Backups	Up to 90 days baseline unless infrastructure or contract defines another approved cycle.	Backup rotation/expiry.	Legal hold or disaster-recovery requirement.

Data type	Baseline retention	Deletion trigger	Possible exception
Data request records	3 years baseline after closure, storing metadata and evidence rather than unnecessary content.	Record expiry.	Legal or dispute requirement.
Privacy / AI incident records	5 years baseline or longer where required.	Record expiry.	Legal, regulatory or contractual requirement.

Baseline periods are internal defaults, not universal legal deadlines. The system-specific schedule must be reconciled with applicable law, client contracts and provider capabilities.

Appendix D. Deletion Propagation Checklist

#	Deletion step
1	Validate requester, authority and deletion scope.
2	Identify client, tenant, user, source document/record IDs and relevant date range.
3	Check legal hold, regulatory retention and contractual exceptions.
4	Disable active access and retrieval where immediate containment is required.
5	Delete or tombstone source file and application record.
6	Delete extracted text, OCR output and temporary parser artifacts.
7	Delete or disable document chunks.
8	Delete or disable embeddings and vector metadata.
9	Invalidate search indexes, caches and conversation memory linked to the source.
10	Delete AI outputs, summaries or workflow artifacts where deletion applies.
11	Evaluate prompt/output content logs and support copies.
12	Submit provider-side deletion where persistent provider data exists and the service supports it.
13	Record backup expiry or deletion-replay treatment.
14	Verify active systems no longer retrieve or display the deleted data.
15	Record completion, exceptions, systems checked, timestamps and responsible owner.

Deletion evidence

The deletion record should prove the process without becoming a new copy of the deleted content. Prefer source identifiers, system names, status codes and timestamps over reproducing personal data in the deletion log.

Appendix E. Data Request Workflow

Stage	Required action	Primary owner
Receive	Record date/time, channel, requester and requested right or client action.	Privacy / Support
Verify	Confirm identity and authority using proportionate means.	Privacy / Client owner
Determine role	Identify whether KOMGIS is controller or processor for the affected processing.	Privacy / Legal
Acknowledge / route	Acknowledge where appropriate; route processor requests to client/controller without delay.	Privacy / Support
Scope systems	Use Data Processing Map to identify source, prompts, outputs, embeddings, providers, logs and backups.	Technical owner
Search / preserve status	Locate data deterministically; apply restriction flag if the request requires processing restriction.	Technical owner
Review exceptions	Assess legal hold, rights exceptions, third-party data and disclosure risk.	Privacy / Legal
Execute	Export, correct, restrict or delete data and propagate changes to derived AI data.	Technical owner
Verify	Check completeness and active retrieval behavior; reconcile provider and backup status.	Privacy / Technical owner
Respond	Provide accurate, understandable response within applicable deadline.	Privacy / Client owner

Stage	Required action	Primary owner
Close	Record outcome, dates, systems, exceptions and evidence; trigger corrective action if a control failed.	Privacy

Appendix F. Provider Privacy Review Checklist

#	Review item	Complete / Evidence
1	Exact provider, product/API, account type and model family are identified.	☐ _____
2	The data categories and maximum sensitivity that may be sent are documented.	☐ _____
3	Provider role and DPA / contractual processing terms are reviewed.	☐ _____
4	Training, fine-tuning, service-improvement and human-review use of customer content are understood.	☐ _____
5	Default and configurable input/output retention are documented.	☐ _____
6	Abuse-monitoring or safety retention that cannot be disabled is documented.	☐ _____
7	Processing and storage regions are identified.	☐ _____
8	Data residency requirements are supported and correctly configured.	☐ _____
9	Applicable international transfer mechanism is assessed.	☐ _____
10	Subprocessor list and change process are reviewed.	☐ _____
11	Provider security documentation and incident commitments are reviewed.	☐ _____
12	Deletion and data-request support capabilities are understood.	☐ _____
13	Enterprise no-training, zero-data-retention or reduced-retention settings are enabled where required and available.	☐ _____
14	Client approval or contractual notification requirements are satisfied.	☐ _____
15	Provider account access and API keys are least-privilege and centrally managed.	☐ _____
16	A provider change / exit plan exists for material client services.	☐ _____
17	The Vendor/Subprocessor List and Data Processing Map are updated.	☐ _____
18	Technical configuration has been verified against the approved provider review.	☐ _____

Appendix G. Required Privacy Records

Record	When required	Owner
Data Processing Map	Required for material AI/client-data systems.	System owner / Technical owner
Controller/processor role assessment	Required where role is not obvious.	Privacy / Legal
Client DPA / processing instructions	Required where KOMGIS acts as processor.	Client owner / Legal
Record of processing activities	Where applicable under law and KOMGIS governance.	Privacy
Data classification	Required for material systems.	System owner
Retention schedule	Required for material systems.	System owner / Privacy
Provider privacy review	Required before sending confidential/personal data to a provider.	Privacy / Security / Technical owner
Vendor/Subprocessor List	Required for approved third parties processing client data.	Privacy / Procurement
Transfer / residency assessment	Required where cross-region or international processing is relevant.	Privacy / Legal
DPIA / privacy impact assessment	Required where screening identifies high risk or law requires it.	Privacy / System owner
Data request register	Required for requests received by KOMGIS.	Privacy / Support
Deletion / offboarding record	Required for material client deletion or service termination.	Technical owner / Client owner

Record	When required	Owner
Privacy incident record	Required for reportable or material privacy events.	Incident owner / Privacy
Approved exception / legal hold	Required whenever policy deletion or processing rules are overridden.	Privacy / Legal / Owner
Training and access evidence	Required according to role and system sensitivity.	Management / Security

Appendix H. Reference Documents

- KOMGIS Responsible AI and AI Management Policy - company-wide AI governance, lifecycle, risk, oversight and accountability.
- KOMGIS AI Security Policy - access control, prompt injection, RAG security, agent tools, secrets, output security and secure development.
- KOMGIS AI Incident Response Plan - incident severity, containment, evidence, investigation, communication, corrective action and reactivation.
- Regulation (EU) 2016/679 - General Data Protection Regulation (GDPR).
- Law on Personal Data Protection, Official Gazette of the Republic of Serbia No. 87/2018.
- European Data Protection Board Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models.
- Regulation (EU) 2024/1689 - Artificial Intelligence Act, including privacy and data-governance considerations relevant to AI systems.