

KOMGIS

Responsible AI and AI Management Policy

Company-wide policy for governing, building, deploying and operating AI systems

Core policy commitment

KOMGIS will use artificial intelligence to improve business processes without weakening privacy, security, human accountability or client trust. AI systems must be governed throughout their lifecycle and must not bypass access control, client data isolation, human approval or documented risk controls.

Policy owner: AI Governance / Technical Leadership

Version: 1.0

Effective date: 8 July 2026

Classification: Internal Policy

Document Control

Field	Value
Document title	KOMGIS Responsible AI and AI Management Policy
Document owner	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 systems, client implementations and approved third parties

Status and legal note

This policy is an internal governance framework informed by KOMGIS's AI security, data privacy and incident-response documents, the management-system structure of ISO/IEC 42001:2023, and Regulation (EU) 2024/1689 (EU AI Act). Adoption of this policy does not by itself constitute ISO certification, a conformity assessment, or a legal determination that a specific AI system meets every regulatory obligation. Each system must be assessed for its actual role, intended purpose, jurisdiction and risk classification.

Policy Statement

KOMGIS designs and operates AI as a controlled component of a business process. AI may assist users, retrieve authorized information, generate content, classify or extract data, and support workflow execution. Responsibility for the system, its intended use, its controls and its business outcome remains with KOMGIS and the authorized human owners of the process.

KOMGIS will establish documented governance for AI systems, identify and treat AI risks, protect client and personal data, apply secure-by-default engineering, provide effective human oversight for sensitive uses, monitor deployed systems, respond to incidents, and continually improve the AI management system.

The rigor of controls shall be proportionate to the system's risk, autonomy, data sensitivity, affected persons and potential impact. Regulatory classification under applicable law shall be evaluated separately from KOMGIS's internal operational risk classification.

Contents

1. Purpose
2. Scope
3. Governance Basis and Policy Hierarchy
4. Definitions
5. Context and AI Management System
6. Leadership, Roles and Accountability
7. Responsible AI Principles
8. AI System Inventory and Intended Purpose
9. AI Risk Classification
10. Prohibited and Restricted AI Uses
11. AI Risk and Impact Assessment
12. Lifecycle Governance and Change Control
13. Data Governance and Privacy
14. Access Control and Client Data Isolation
15. Model Provider, Vendor and Subprocessor Governance
16. RAG and Document Intelligence Security
17. Prompt Injection and Untrusted Content
18. AI Agents, Tools and Automation
19. Human Oversight and Approval
20. Transparency, Traceability and Source Citations
21. Accuracy, Evaluation and Testing
22. Secure Development, Deployment and Operations
23. Logging, Auditability and Monitoring
24. Retention, Deletion and Backups
25. AI Literacy, Competence and Awareness
26. AI Incident Response
27. Client Rights, Requests and Communication
28. Nonconformity and Corrective Action
29. Internal Audit, Management Review and Improvement
30. Policy Review and Enforcement
- Appendix A. Internal AI Risk Classification Matrix
- Appendix B. Required AI Governance Records
- Appendix C. AI Incident Severity Matrix
- Appendix D. Pre-Deployment AI Review Checklist
- Appendix E. Alignment Overview

1. Purpose

This policy defines how KOMGIS governs artificial intelligence systems and AI-enabled software throughout planning, design, development, testing, deployment, use, monitoring, change and retirement.

The policy is intended to ensure that AI is:

- used for an approved and documented business purpose;
- human-centric and subject to clear accountability;
- proportionately governed according to risk and impact;
- secure against unauthorized access, prompt injection, data leakage and unsafe automation;
- privacy-aware and limited to necessary data processing;
- traceable, testable and monitored throughout its operational lifecycle;
- supported by competent personnel and effective human oversight;
- managed through documented incident, corrective-action and continual-improvement processes.

The minimum internal rule is that an AI model, prompt or agent is never a security or accountability boundary. Security and control must be enforced through application logic, permissions, infrastructure, validation, monitoring and human oversight.

2. Scope

This policy applies to AI systems developed, configured, integrated, deployed or used by KOMGIS, including:

- AI assistants and chatbots;
- retrieval-augmented generation (RAG) and enterprise document intelligence platforms;
- document summarization, extraction, classification and search systems;
- AI agents and tool-using systems;
- voice interfaces, conversational interfaces and AI-enabled avatars;
- workflow automation and decision-support systems;
- AI-generated reports, drafts and structured outputs;
- AI coding assistants used for company work;
- embeddings, vector databases, model gateways and model integrations;
- client-specific AI systems and internal AI tools.

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

It applies across development, testing, staging, production, client deployments, cloud infrastructure, APIs, databases, file stores, vector stores, logs, monitoring, backups and third-party integrations.

3. Governance Basis and Policy Hierarchy

KOMGIS's AI governance framework is structured around management-system practices for context, leadership, planning, support, operation, performance evaluation and improvement. This reflects the clause structure visible in ISO/IEC 42001:2023 and is adapted to KOMGIS's organization and AI delivery model.

For systems within the scope of the EU AI Act, KOMGIS shall determine its role for the specific system, such as provider or deployer, and assess the system's regulatory classification and applicable obligations before placing the system on the market, putting it into service or using it.

This policy is supported by more specific documents. Where a specific policy sets a stricter requirement, the stricter requirement applies.

Supporting document	Primary subject
AI Security Policy	Access control, prompt injection, RAG, agent tools, secrets, output security and secure development
Data Privacy Policy	Client data, AI inputs and outputs, embeddings, providers, retention, deletion and data requests
AI Incident Response Plan	AI incident severity, containment, investigation, communication, corrective

Supporting document	Primary subject
	action and reactivation
AI System Inventory	System ownership, intended use, models, data, risk and operational status
AI Risk Assessment / Impact Assessment	Known and foreseeable risks, affected parties, controls and residual risk
Model Provider Disclosure / Vendor List	Third-party AI and cloud processing
Human Oversight Record	Review and approval of sensitive AI outputs or actions
Evaluation and Testing Records	Pre-deployment and ongoing quality, security and safety evidence

4. Definitions

Term	Policy meaning
AI system	A machine-based system that infers from inputs to generate outputs such as content, predictions, recommendations or decisions that may influence physical or virtual environments.
AI system owner	The accountable business owner for an AI system's intended purpose, risk, approval, operation and retirement.
Technical owner	The person or team accountable for the system's technical implementation, controls, testing and operational integrity.
Provider	A role that develops an AI system or has it developed and places it on the market or puts it into service under its name or trademark, as assessed under applicable law.
Deployer	A role that uses an AI system under its authority in a professional or organizational context, as assessed under applicable law.
Client data	Information provided by, generated for, or connected to a client, including prompts, outputs, embeddings and workflow records.
RAG	A system that retrieves information from documents, databases or knowledge stores and provides selected context to an AI model.
AI agent	An AI-enabled system able to select or call approved tools, integrations or actions to advance an assigned task.
Sensitive action	An external, irreversible, high-impact or regulated action that requires defined human approval.
AI incident	An event involving harmful, wrong, unsafe, unauthorized or policy-violating AI behavior or output that creates actual or potential impact.
Residual risk	Risk remaining after defined controls and mitigations are applied.

5. Context and AI Management System

KOMGIS shall maintain an AI management system appropriate to its role as a developer and operator of AI-enabled business software. The system shall consider KOMGIS's business objectives, client expectations, contractual commitments, data-processing activities, technical architecture, third-party providers and applicable legal or regulatory requirements.

The AI governance scope shall be reviewed when KOMGIS:

- enters a new industry or regulated domain;
- introduces a materially new AI capability, such as autonomous tool use or biometric processing;
- adds a new model provider, hosting region or significant subprocessor;
- changes from decision support to decision automation;

- begins processing a materially different category of client or personal data;
- changes a system's intended purpose, user group or deployment context.

Interested parties may include clients, end users, employees, contractors, affected persons, model and infrastructure providers, regulators and other parties with contractual or legal rights. Relevant requirements shall be captured in system or project documentation.

6. Leadership, Roles and Accountability

Leadership is accountable for establishing the direction of responsible AI use, providing appropriate resources, assigning roles, supporting reporting of concerns and ensuring that business pressure does not override required AI controls.

Role	Minimum responsibilities
Executive leadership	Approve policy direction; ensure resources; review high or critical risk systems and material incidents; support corrective action.
AI governance / technical leadership	Maintain the AI governance framework; approve risk methodology; review model providers; coordinate audits and management review.
AI system owner	Define intended use; maintain business accountability; approve risk acceptance; ensure human oversight and operational review.
Technical owner	Implement architecture and controls; maintain testing, logging, monitoring, change control and technical incident support.
Security / privacy owner	Review data exposure, access control, prompt injection, client isolation, provider risk and privacy requirements.
Project manager / client owner	Capture client requirements; ensure project governance records are complete; coordinate approved client communication.
Human reviewer	Understand AI limitations; inspect relevant source evidence; approve, reject, correct or escalate sensitive outputs and actions.
Employees and contractors	Use approved AI tools; protect data; verify outputs; follow approval controls; promptly report unsafe or suspicious AI behavior.

Each production AI system shall have a named AI system owner and technical owner. Responsibility shall not be assigned solely to the model provider.

7. Responsible AI Principles

7.1 Human accountability

AI assists a controlled business process. An accountable natural person or organizational role remains responsible for the system and for sensitive business outcomes.

7.2 Risk proportionality

Controls are increased as autonomy, data sensitivity, potential harm, affected population, scale and regulatory exposure increase.

7.3 Purpose limitation

An AI system shall be used only for a documented intended purpose and approved reasonably foreseeable uses.

7.4 Data minimization

AI receives only the data needed for the approved task. Entire document libraries or databases shall not be sent when limited excerpts or records are sufficient.

7.5 Least privilege

Users, agents, tools, APIs and service accounts receive the minimum access and action scopes needed.

7.6 Client isolation

One client's users, AI systems, retrieval processes, agents and credentials shall not access another client's data.

7.7 Transparency and traceability

AI involvement, source material, model or workflow versions, approvals and significant actions shall be traceable according to risk.

7.8 Security by design

AI-specific risks are addressed through application and infrastructure controls, not prompt wording alone.

7.9 Human oversight

Humans must be able to understand relevant limitations, identify anomalous behavior, disregard or override AI output, and stop sensitive automation.

7.10 Continual improvement

Monitoring, incidents, evaluation results, complaints and audit findings shall feed corrective and preventive action.

8. AI System Inventory and Intended Purpose

No production AI system shall be operated for KOMGIS or a client unless it is recorded in the AI System Inventory.

The inventory shall identify, at minimum:

- system name, version and status;
- AI system owner and technical owner;
- intended purpose and expected users;
- KOMGIS's assessed role under applicable law where relevant;
- internal risk level and regulatory classification assessment;
- models and model providers used;
- data categories processed, including personal or confidential data;
- RAG data sources, vector stores and major integrations;
- available tools and whether they are read or write actions;
- human oversight and approval requirements;
- deployment environment, hosting region and relevant subprocessors;
- evaluation, monitoring and incident links;
- retention and deletion rules.

The intended purpose shall be specific enough to evaluate performance and risk. Statements such as 'general AI assistant' are insufficient for systems that access client data or perform actions.

9. AI Risk Classification

Every AI system shall receive an internal KOMGIS operational risk classification before production use. The internal classification is a governance mechanism and does not replace regulatory classification under the EU AI Act or other law.

Level	Typical characteristics	Minimum governance
Low	Internal assistance; non-sensitive data; no external or material action; errors are easily reversible.	Inventory, owner, acceptable-use controls, basic testing and monitoring.
Limited	Client-facing content or business workflow support; confidential data; RAG; structured extraction; limited integrations.	Risk assessment, privacy/security review, documented testing, logs, monitoring and defined human review.
High	Sensitive domain, important decisions, material external	Formal risk and impact assessment, strong

Level	Typical characteristics	Minimum governance
	communication, write tools, significant personal data, high autonomy or potentially serious client impact.	oversight, formal approval gates, audit logs, security testing, legal/compliance review and management risk acceptance.
Critical	Potential severe harm, large-scale cross-client exposure, dangerous autonomous action, safety impact or use that may fall within prohibited or high-risk regulatory categories.	Do not deploy until executive, technical, security and legal/compliance approval; formal control plan and enhanced monitoring. Prohibited uses remain prohibited.

Risk classification shall consider:

- intended use and reasonably foreseeable misuse;
- severity and likelihood of harm to individuals, clients or KOMGIS;
- whether the system influences legal, financial, health, employment, education, insurance, eligibility or safety outcomes;
- data sensitivity, scale and client isolation risk;
- autonomy, tool permissions and reversibility of actions;
- reliance on AI output and risk of automation bias;
- accuracy requirements and consequences of hallucination or extraction error;
- prompt injection, adversarial manipulation and cybersecurity exposure;
- affected persons, including minors or vulnerable groups;
- regulatory classification and contractual obligations.

10. Prohibited and Restricted AI Uses

KOMGIS shall not approve AI uses that are unlawful or conflict with this policy. In addition, KOMGIS adopts a conservative internal position toward practices identified as prohibited by the EU AI Act.

The following uses are prohibited as ordinary KOMGIS commercial or internal AI work:

- AI designed to use subliminal, purposefully manipulative or deceptive techniques to materially distort behavior in a manner likely to cause significant harm;
- AI designed to exploit vulnerabilities arising from age, disability or social or economic situation to materially distort behavior and cause or likely cause significant harm;
- social scoring that leads to unjustified, disproportionate or context-unrelated detrimental treatment;
- assessment of a person's criminal-offence risk based solely on profiling or personality characteristics;
- untargeted scraping of facial images from the internet or CCTV to create or expand facial-recognition databases;
- emotion inference in workplace or education contexts, except a use that is specifically assessed as lawful for medical or safety reasons;
- biometric categorization intended to deduce or infer highly sensitive traits where prohibited by applicable law;
- real-time remote biometric identification in publicly accessible spaces for law-enforcement purposes, except where a competent legal review determines a specific lawful exception applies.

Restricted use rule

Any project involving biometrics, emotion recognition, vulnerable persons, employment, education, health, insurance, credit, essential services, law enforcement, migration, justice, democratic processes or safety components must be escalated before development commitment. The project may proceed only after regulatory classification, risk assessment and written approval.

11. AI Risk and Impact Assessment

Risk management shall be a documented, iterative process carried out throughout the lifecycle. Risk assessment is required before launch for limited, high and critical systems and whenever a material change affects risk.

The assessment shall:

1. identify known and reasonably foreseeable risks under intended use;
2. consider reasonably foreseeable misuse and abnormal operating conditions;

3. identify affected persons, clients, business processes, data and systems;
4. estimate likelihood and impact using the approved KOMGIS methodology;
5. identify existing and planned preventive, detective and corrective controls;
6. evaluate residual risk after controls;
7. define testing and monitoring needed to verify the controls;
8. assign owners and due dates for treatment actions;
9. record risk acceptance by an authorized role;
10. use post-deployment evidence, incidents and monitoring to update the risk assessment.

An AI system impact assessment shall be performed where the potential impact on people, fundamental rights, safety or important decisions is material. It should consider:

- intended benefits and necessity of using AI;
- affected individuals or groups and potential uneven impact;
- data provenance, representativeness and bias risks where relevant;
- explainability and source evidence needed for users or reviewers;
- human oversight competence, authority and workload;
- ability to contest, correct, override or reverse the output or action;
- privacy, confidentiality, security and cross-client exposure;
- failure scenarios and incident response capability.

12. Lifecycle Governance and Change Control

AI governance applies to the full lifecycle:

Lifecycle stage	Required governance focus
Concept / sales	Define use case, intended purpose, client requirements, data categories and obvious prohibited or high-risk concerns.
Design	Complete system inventory entry, architecture, risk classification, data-processing map, provider review and oversight design.
Build	Secure development, prompt/workflow versioning, access control, retrieval filtering, tool scoping and documented technical controls.
Test	Evaluate quality, security, permissions, prompt injection, RAG retrieval, sensitive actions, failure handling and defined performance criteria.
Approve / deploy	Resolve blocking risks, approve residual risk, document release, confirm monitoring, incident route and rollback/disable mechanism.
Operate	Monitor performance, incidents, complaints, tool calls, retrieval sources, approval rates and material drift.
Change	Assess impact of model, prompt, provider, data source, tool, permission, intended-purpose or workflow changes before release.
Retire	Disable access and automation, execute retention/deletion rules, revoke credentials and preserve required records.

Material changes require formal review. Material changes include:

- new intended purpose or user population;
- new model provider or materially different model family;
- new client data category, data source or processing region;
- new write tool, integration or external action;
- increase in autonomy or removal of a human approval step;
- material prompt or workflow behavior change;
- new RAG source or change to permission filtering;
- significant performance degradation, incident pattern or known vulnerability.

Prompt, workflow and system changes shall be versioned so that KOMGIS can identify the version associated with a material output or incident when required.

13. Data Governance and Privacy

Client data includes account data, documents, business and workflow records, prompts, retrieved context, AI outputs, embeddings, vector metadata, audit logs, approvals and support records. Derived AI data shall be governed as client data when it is generated from client content.

AI data processing shall follow these rules:

- process data only for an approved, documented business purpose;
- collect and send only data necessary for the AI task;
- use approved systems and secure channels;
- do not sell client data or use it for unrelated advertising;
- do not use client data to train or fine-tune company or third-party models unless the client explicitly approves the use and the purpose, data, legal/contractual requirements, retention and safeguards are documented;
- treat prompts, outputs, embeddings and vector data as governed data;
- maintain a Data Processing Map for material AI systems;
- apply privacy by design and by default where personal data is processed;
- define retention and deletion for source and derived AI data.

Development and testing shall use synthetic, dummy, anonymized or client-approved test data whenever possible. Real production client data may be used only when necessary, approved, restricted, documented and removed when no longer needed.

International or cross-region processing shall be documented where relevant. Client data residency requirements must be captured before deployment and reflected in the system architecture and provider configuration.

14. Access Control and Client Data Isolation

Authentication and authorization are foundational AI controls. Frontend hiding is not authorization. Access must be enforced in backend application logic and at the relevant data, file, retrieval and tool layers.

Controls shall include, according to system risk:

- authenticated access unless the feature is explicitly public;
- role-based access control and object-level authorization;
- tenant, organization, workspace or project identifiers on client-owned records;
- tenant-aware database queries and file access checks;
- client- and user-scoped vector retrieval filters applied before context enters the model prompt;
- client-scoped integrations and credentials;
- least-privilege service accounts and API keys;
- periodic access review for sensitive systems;
- time-bound and logged production-data access for developers or support personnel where practical.

Mandatory isolation rule

The system must retrieve only data the active user and client workspace are authorized to access. It is prohibited to retrieve broadly and rely on the model to hide unauthorized information afterward.

15. Model Provider, Vendor and Subprocessor Governance

Client or confidential data may be sent only to approved model providers and subprocessors for the approved use case.

Before approval, KOMGIS shall review, as relevant:

- provider, product and model family;
- purpose of use and data sent to the provider;
- whether personal, confidential or regulated data may be processed;
- training-use and service-improvement terms;

- retention and deletion terms;
- security documentation and available enterprise controls;
- processing region and data-residency options;
- contract, DPA and subprocessor status;
- client contractual or approval requirements;
- service continuity and provider-change risk.

KOMGIS shall maintain a Model Provider Disclosure and Vendor/Subprocessor List appropriate to its services. Personal AI accounts, unknown providers and unapproved public AI tools shall not be used for client confidential data.

A provider or model change is a controlled change and shall trigger review of the AI System Inventory, Data Processing Map, vendor records, risk assessment, retention settings, client requirements, evaluation evidence and security controls.

16. RAG and Document Intelligence Security

RAG and document intelligence systems are core KOMGIS use cases and shall be designed so that source authorization is enforced before retrieval. Source citations and traceability are required where the system is intended to answer from enterprise documents.

Every production RAG system shall implement, where applicable:

1. client- or tenant-level filtering;
2. user-level permission checks;
3. document- or object-level authorization;
4. metadata filtering in vector or hybrid search;
5. source and chunk identifiers sufficient for traceability;
6. permission-aware prompt construction;
7. secure vector database access and client-scoped metadata;
8. treatment of retrieved content as untrusted data;
9. controls for deletion or disabling of chunks and embeddings when source data is deleted;
10. retrieval evaluation and cross-client isolation testing.

For KOMGIS's enterprise document intelligence platforms, the authorization hierarchy used by the application - for example Project → Field → Section → Document or another client-defined hierarchy - must also be reflected in retrieval permissions. A model must not gain broader document access than the authenticated user.

Where the intended answer is grounded in source material, the user interface should display inline citations or source references that allow the user or reviewer to identify the supporting document and relevant source location.

17. Prompt Injection and Untrusted Content

User input, uploaded files, webpages, emails, support tickets, contracts, spreadsheets, retrieved documents and tool outputs shall be treated as untrusted content. Instructions found inside such content shall not override system policy, authorization, tenant isolation, tool permissions or human approval requirements.

Prompt injection controls shall include defense in depth:

- clear separation of trusted system instructions from external data;
- permission filtering before prompt construction;
- explicit tool allowlists and server-side permission checks;
- separation of read and write tools;
- validation of sensitive outputs and structured schemas;
- limits on data export and bulk retrieval;
- prompt injection and malicious-document test cases;
- monitoring of suspicious retrieval, tool or export patterns;
- safe handling of hidden text, PDF layers, comments, OCR artifacts and other document content.

Prompt wording alone is not an adequate security control. The model shall not decide its own permissions, tenant, tool scopes or whether a human approval requirement applies.

18. AI Agents, Tools and Automation

AI agents may use only explicitly approved tools for the active use case, user and client workspace. Agents shall be designed as controlled workflow participants rather than unrestricted actors.

Agent controls shall include:

- documented tool inventory and purpose;
- least-privilege tool and credential scopes;
- separate read and write capabilities where practical;
- server-side authorization before every sensitive tool call;
- human approval before sensitive actions;
- action preview or clear confirmation for material external actions where appropriate;
- logging of tool calls, action status and approval status;
- rate, volume and export controls where relevant;
- a shutdown, disable or safe-stop mechanism;
- ability to revoke a tool or credential without disabling unrelated systems.

Examples of sensitive write actions that require human approval include:

- sending external email or customer communication;
- updating official client, legal, HR, insurance or financial records;
- triggering a payment or financial change;
- sharing files or data externally;
- submitting forms or approvals;
- deleting important records;
- closing a complaint or case where the outcome affects a person or client;
- making another material or difficult-to-reverse change.

An AI agent shall not bypass approval by retrying, rephrasing a requested action, selecting a different tool, or splitting one sensitive action into multiple smaller calls.

19. Human Oversight and Approval

Human oversight shall be appropriate to the system's risk, autonomy and context. Reviewers must have the competence, training, authority and access to evidence needed to perform meaningful oversight.

For systems requiring human oversight, reviewers shall be able to:

- understand the relevant capabilities and limitations of the AI system;
- recognize automation bias and avoid relying on AI output merely because it is confident or well written;
- inspect sources, retrieved documents or other relevant evidence;
- correctly interpret the output in the context of the business process;
- reject, disregard, override or reverse the output when appropriate;
- stop or disable the affected workflow when anomalous or unsafe behavior is detected;
- escalate uncertainty or sensitive cases.

Where approval is required, the system shall block the sensitive action until approval is received from an authorized person. Approval records should capture the reviewer, output reviewed, material edits, approved action, decision, timestamp and related case or workflow identifier.

AI recommendations shall not be treated as final decisions in sensitive workflows unless a specific lawful automation use has been assessed, approved and documented.

20. Transparency, Traceability and Source Citations

KOMGIS shall provide clear information about AI involvement where required by law, contract or the nature of the user interaction. Systems intended to interact directly with natural persons shall inform users that they are interacting with an AI system when this is not otherwise obvious and when the applicable transparency obligation applies.

Where KOMGIS provides or deploys systems generating synthetic content, applicable marking or disclosure obligations shall be assessed. Client-facing AI transparency notices shall describe the AI-assisted nature of the feature, the purpose of processing, relevant limitations and the role of human review where appropriate.

Traceability records may include:

- AI system and model identifier;
- prompt or workflow version;
- request timestamp and user or service identity;
- client or workspace identifier;
- retrieved source identifiers;
- tool calls and action status;
- approval or rejection status;
- final action or output disposition;
- evaluation or incident reference where applicable.

For document-grounded answers, citations should point back to the authorized source material. Citations support verification but do not replace permission checks, quality evaluation or human review.

21. Accuracy, Evaluation and Testing

AI systems must be evaluated against their intended purpose before production use. Testing shall use defined metrics, acceptance criteria and representative scenarios appropriate to the system and its risk.

Evaluation may include:

- task accuracy and completeness;
- faithfulness to source documents and hallucination rate;
- citation correctness and source coverage;
- field extraction precision and validation failures;
- classification performance and error distribution;
- human edit, rejection and escalation rates;
- latency, availability and cost limits;
- bias or uneven performance where relevant to affected persons;
- robustness to malformed, ambiguous or adversarial inputs;
- failure behavior and ability to fall back to manual processing.

Required security and control tests shall be selected according to the architecture and risk. They may include:

Test	Purpose
Authentication	Confirm only valid users or systems can access the service.
Authorization / object access	Confirm users can access only approved resources.
Tenant isolation	Confirm clients cannot access one another's data.
RAG retrieval permissions	Confirm only authorized documents and chunks are retrieved.
Prompt injection	Test malicious user and indirect-document instructions.
Agent tool permissions	Confirm agents cannot invoke unauthorized tools or scopes.
Sensitive action approval	Confirm write actions are blocked until required approval.
Output validation	Confirm schemas, allowed fields and downstream safety checks.
Log redaction	Confirm secrets and unnecessary sensitive content are not over-logged.
File upload	Confirm type, size, storage and safe-processing controls.
Deletion	Confirm source and derived data are removed or disabled according to policy.
Rate and abuse controls	Confirm defined volume, export and abuse limits.

High-risk or critical internal systems require formal test evidence and approval before deployment. Known blocking defects, failed isolation tests or bypassed approval gates prohibit production release.

22. Secure Development, Deployment and Operations

AI-enabled software shall follow secure development and deployment practices. Developers shall:

- perform code review and appropriate testing;
- enforce backend authorization and tenant isolation;
- validate inputs and outputs;
- protect secrets and avoid hardcoded credentials;
- use approved dependency and vulnerability management practices;
- secure file upload and extraction pipelines;
- review model and provider configuration;
- test RAG permissions, prompt injection and agent tools;
- avoid real client data in development unless approved;
- maintain rollback or disable capability for material AI features.

AI-generated code, commands or infrastructure changes must be reviewed by a competent developer before use. AI output must not be executed directly as SQL, shell commands, scripts or code except inside a specifically designed and approved controlled execution environment.

Secrets such as passwords, API keys, access tokens, private keys, database credentials and encryption keys must not be placed in AI prompts or exposed to models.

If a secret is accidentally exposed, the event must be reported, the secret revoked or rotated, relevant logs reviewed, provider exposure assessed and preventive controls updated.

23. Logging, Auditability and Monitoring

AI systems shall be monitored for security, reliability, abuse, operational risk and material performance degradation.

Logging shall be proportionate to the risk and shall minimize unnecessary sensitive content.

Depending on risk, logs may record:

- user or service identity and organization/workspace identifier;
- request timestamp and AI system identifier;
- model and prompt/workflow version;
- retrieved source identifiers;
- tool calls and action results;
- human approval status;
- error, validation or security status;
- latency, token usage or cost metadata;
- incident or alert correlation identifiers.

Logging controls shall:

- prefer metadata over full prompt or document content when possible;
- redact personal data and secrets where practical;
- restrict access to full prompt and output logs;
- separate audit and debug logs where appropriate;
- apply retention limits;
- protect audit logs against unauthorized modification.

Monitoring signals may include:

- high reviewer rejection or edit rates;
- repeated hallucinations or unsupported claims;
- unexpected retrieval sources or wrong-client references;
- prompt injection attempts;
- unauthorized or unusual tool calls;
- unexpected agent actions or approval bypass attempts;
- large extraction, export or retrieval volume;
- sensitive data detected in outputs or logs;
- abnormal token usage or cost spikes;
- client complaints and repeated support issues.

For high-risk systems subject to specific legal requirements, post-deployment monitoring and serious-incident reporting obligations shall be assessed and implemented according to the applicable role and regulatory classification.

24. Retention, Deletion and Backups

Client and AI data shall be retained only as long as needed for the approved purpose, contract, legal requirement, security need or documented audit requirement.

Data type	Policy expectation
Uploaded source files	Client- or service-defined retention; protected and deleted according to contract and policy.
Extracted text and document chunks	Normally linked to the source document's lifecycle.
Embeddings / vector records	Governed as derived client data; normally deleted or disabled with the source unless a documented exception applies.
AI outputs / chat history	Client- or workflow-defined retention proportionate to business and audit needs.
Audit logs	Longer retention may be required for security, traceability or compliance.
Debug / error logs	Short technical retention unless needed for an active incident.
Backups	Retained according to the backup schedule and protected against unauthorized access.

Deletion workflows shall consider source and derived data, including:

- uploaded files and database records;
- extracted text and cached content;
- document chunks and embeddings;
- AI outputs and reports;
- chat history and temporary files;
- integration data and caches;
- logs where deletion is applicable;
- backups according to the backup-retention process.

When a client document is deleted, associated active extracted text, chunks, embeddings and AI-generated summaries should be deleted or disabled according to defined retention and audit requirements.

25. AI Literacy, Competence and Awareness

KOMGIS shall take measures to maintain a sufficient level of AI literacy for personnel who develop, operate, supervise or use AI systems on KOMGIS's behalf. Training shall consider the person's technical knowledge, experience, role, system context and affected users.

Audience	Minimum competency themes
All staff using AI	Approved tools, confidentiality, data handling, hallucinations, verification, prompt injection awareness, incident reporting.
Developers	AI architecture, authorization, RAG permissions, secure prompting, tool scoping, model/provider controls, evaluation, monitoring and incident support.
Human reviewers	System capabilities and limits, source verification, automation bias, approval criteria, rejection and escalation.
Managers / system owners	Risk classification, intended purpose, oversight, residual-risk acceptance, change control and client requirements.
Support / administrators	Least privilege, client data handling, access logging, troubleshooting boundaries and incident preservation.

Training completion and role-specific competence should be documented for personnel assigned to high-impact AI operation or oversight. AI literacy needs shall be reviewed when systems, tools or regulatory obligations materially change.

26. AI Incident Response

An AI incident is any event in which an AI system produces, supports, sends, exposes, recommends or performs something that may cause harm, risk, policy violation, client impact or loss of trust.

Examples include:

- wrong or misleading output used in a business process;
- dangerous or harmful recommendation;
- confidential or personal data exposure;
- wrong-client retrieval or cross-client data leakage;
- successful prompt injection or manipulation;
- unauthorized AI agent action or tool call;
- sensitive external communication sent without required review;
- bias or unfair treatment with material impact;
- failure of required human oversight;
- security, privacy, contractual or regulatory failure involving AI.

When an AI incident is suspected, KOMGIS shall:

1. stop use of the affected output and contain ongoing impact;
2. disable the affected feature, agent or tool when necessary;
3. preserve relevant prompts, source references, logs, approvals and other evidence;
4. classify severity and assign an incident response owner;
5. identify affected systems, data, clients, users and actions;
6. revoke or rotate exposed credentials where needed;
7. investigate the technical, process and human causes;
8. coordinate client, legal, compliance or regulatory review where required;
9. correct affected outputs or actions and implement root-cause controls;
10. retest relevant controls before reactivation;
11. document the incident, residual risks and corrective actions;
12. feed lessons learned into risk, testing, monitoring and policy updates.

Employees must immediately report clearly wrong, harmful, unauthorized or suspicious AI behavior. Serious incidents must not be hidden, silently corrected without reporting, or dismissed on the basis that AI systems sometimes make mistakes.

Reactivation rule

A disabled AI feature shall not be reactivated until containment is confirmed, root cause is understood to an appropriate level, corrective controls are implemented and tested, required notifications or reviews are addressed, and the authorized system owner approves reactivation. High or critical incidents require security/legal/compliance or management approval as defined by the Incident Response Plan.

27. Client Rights, Requests and Communication

KOMGIS shall maintain a process to receive, track and respond to client data and AI-governance requests according to contract and applicable law.

Requests may include:

- access, correction, export, deletion or restriction of client data;
- information about model providers and subprocessors;
- information about data retention or processing regions;
- information about relevant security controls and client data isolation;

- confirmation of KOMGIS's default position that client data is not used for model training without explicit approval;
- information about the AI-assisted nature of a feature and human oversight.

Client incident communication shall be accurate, brief and coordinated. It should state what is known, the affected system or workflow, containment status, known impact, corrective actions and whether the client must take action. KOMGIS shall avoid speculation or blaming a model provider or user without evidence.

Public or regulatory communication concerning material AI incidents must be approved by leadership and legal/compliance as required.

28. Nonconformity and Corrective Action

A nonconformity is a failure to meet this policy, a supporting AI governance requirement, a documented system requirement or an applicable obligation.

Nonconformities may be identified through testing, monitoring, audit, incident review, client complaints, reviewer feedback or management review.

KOMGIS shall:

1. record the nonconformity at a level proportionate to its risk;
2. contain or correct the immediate issue;
3. determine whether similar systems or workflows may be affected;
4. identify root or contributing causes where appropriate;
5. define corrective actions, owners and due dates;
6. verify that corrective actions are effective;
7. update risks, tests, training, monitoring, architecture or documentation where necessary.

Corrective actions may include prompt or workflow changes, RAG permission fixes, access-control changes, provider changes, human approval gates, output validation, tool restriction, additional monitoring, reviewer training, rollback or feature retirement.

29. Internal Audit, Management Review and Improvement

KOMGIS shall periodically evaluate whether the AI management system is implemented and effective. The depth and frequency of review shall be proportionate to KOMGIS's AI portfolio and risk.

Internal audit or governance review should examine a representative sample of:

- AI System Inventory completeness and ownership;
- risk classifications and risk/impact assessments;
- provider and subprocessor approvals;
- data-processing maps, retention and deletion rules;
- RAG authorization and client isolation evidence;
- human oversight and sensitive-action approvals;
- evaluation, testing and release evidence;
- logging, monitoring and alert handling;
- incidents, corrective actions and reactivation decisions;
- AI literacy and role-specific training records.

Management review shall consider:

- changes in KOMGIS's business, clients, legal environment and AI technologies;
- AI objectives and performance indicators;
- material incidents and recurring failure patterns;
- audit findings and corrective-action status;
- model provider, supply-chain and security changes;
- client feedback, complaints and enterprise security-review findings;
- resource, competence and training needs;
- opportunities for improvement and decisions on risk appetite.

Management review outputs shall record key decisions, improvement actions, owners and required changes to the AI governance framework.

30. Policy Review and Enforcement

This policy shall be reviewed at least annually and additionally when:

- a new material AI system or agent capability is introduced;
- a high or critical AI incident occurs;
- a new model provider or important subprocessor is added;
- a material new data category or RAG source is processed;
- human oversight or agent tool permissions materially change;
- client contractual requirements materially change;
- relevant legal, regulatory, security or industry requirements change;
- management review or audit identifies a policy gap.

Violations may result in removal of AI tool access, suspension of a feature or integration, corrective training, project escalation or other action consistent with KOMGIS's contractual and employment arrangements.

No employee or contractor shall be penalized for promptly reporting a good-faith AI safety, security, privacy or compliance concern through approved channels.

Appendix A. Internal AI Risk Classification Matrix

Factor	Low	Limited	High / Critical indicators
Data	Public or non-sensitive	Confidential client data	Sensitive personal/regulated data; large scale; cross-client exposure risk
User impact	Internal convenience	Business workflow support	Material effect on rights, access, money, health, employment, insurance, safety or essential services
Autonomy	No tools / draft only	Read tools or reversible support	Write tools, external action, high autonomy, irreversible or difficult-to-reverse action
Human review	Normal user judgment	Defined review for client-facing or material output	Mandatory approval, trained reviewer, override/stop capability
Accuracy consequence	Minor inconvenience	Rework or client dissatisfaction	Serious financial, legal, safety, privacy or fundamental-rights impact
Attack exposure	Closed internal use	External input or document ingestion	Untrusted RAG sources, public interface, tool access, valuable data or adversarial users
Regulatory signal	No obvious special category	Transparency/privacy obligations possible	Potential prohibited/high-risk AI Act category or other regulated domain

Classification is based on the most consequential reasonably foreseeable scenario, not the average expected use. The owner may classify a system at a higher level when uncertainty is material.

Appendix B. Required AI Governance Records

Record	Low	Limited	High / Critical
AI System Inventory	Required	Required	Required
Named system and technical owners	Required	Required	Required
Risk classification	Required	Required	Required
AI Risk Assessment	Basic / as needed	Required	Formal and approved
AI Impact Assessment	Usually not required	Case-by-case	Required where people or important decisions may be materially affected
Data Processing Map	If client/personal data	Required	Required
Model Provider Disclosure	If provider used	Required	Required
Vendor/Subprocessor List	As applicable	Required	Required
Human Oversight Record	Not normally required	Sensitive actions	Required for defined approvals
Evaluation and Test Record	Basic	Required	Formal
Security Review	Basic	Required	Formal
Monitoring Plan	Basic	Required	Formal / enhanced
Incident Response linkage	Required	Required	Required
Change / release record	Basic	Required	Formal

Appendix C. AI Incident Severity Matrix

Severity	Example impact	Immediate response
1 - Low	Minor formatting or non-sensitive internal error caught before use.	Correct output; record if recurring; adjust prompt or instruction if needed.
2 - Medium	Wrong internal summary; incorrect extraction caught in review; repeated quality failure.	Contain; notify owner; review affected outputs; document; correct workflow; monitor recurrence.
3 - High	Client-facing wrong information; confidential exposure; wrong-client retrieval; unauthorized action; sensitive workflow impact.	Immediate containment; incident owner; preserve logs; identify affected parties; root cause; corrective action; notification review; approval before reactivation.
4 - Critical	Large-scale cross-client leak; severe personal-data exposure; dangerous high-impact automated action; serious harm or major regulatory risk.	Stop affected system; leadership/security/legal response; preserve evidence; urgent impact analysis; required notifications; formal root cause; management approval before reactivation.

Appendix D. Pre-Deployment AI Review Checklist

#	Review item	Complete / Evidence
1	The system is listed in the AI System Inventory and has named owners.	☐ _____
2	The intended purpose, users and deployment context are documented.	☐ _____
3	KOMGIS's legal role and regulatory classification have been assessed where applicable.	☐ _____
4	The internal risk classification is approved.	☐ _____
5	Required AI risk and impact assessments are complete.	☐ _____
6	The Data Processing Map identifies data, transfers, providers and storage.	☐ _____
7	The model provider and subprocessors are approved for the data and use case.	☐ _____
8	Client requirements, data residency and provider restrictions are documented.	☐ _____
9	Authentication, role permissions and object-level authorization are implemented.	☐ _____
10	Tenant and client data isolation are tested.	☐ _____
11	RAG retrieval is filtered by client, user and document permission before prompt construction.	☐ _____
12	Retrieved content and tool outputs are treated as untrusted.	☐ _____

#	Review item	Complete / Evidence
13	Prompt injection defenses and malicious-document tests are complete.	☐ _____
14	Agent tools are explicitly approved and least-privilege scoped.	☐ _____
15	Read and write tools are separated where practical.	☐ _____
16	Sensitive actions are blocked pending human approval.	☐ _____
17	Human reviewers have sufficient evidence, competence, authority and override/stop capability.	☐ _____
18	AI output validation is implemented for downstream APIs or automated workflows.	☐ _____
19	Secrets are excluded from prompts and logs.	☐ _____
20	Required quality, citation, extraction or classification evaluation has passed defined criteria.	☐ _____
21	Security tests applicable to the architecture have passed.	☐ _____
22	Prompt/workflow/model versions are identifiable.	☐ _____
23	Audit logging and sensitive-log controls are configured.	☐ _____
24	Monitoring and alert ownership are defined.	☐ _____
25	Retention and deletion rules cover source and derived AI data.	☐ _____
26	The incident reporting path and disable/rollback mechanism are ready.	☐ _____
27	Required client transparency or AI interaction notice is implemented.	☐ _____
28	Residual risks have been accepted by the authorized owner.	☐ _____

Appendix E. Alignment Overview

The table below is a high-level design alignment. It is not a certification or legal-compliance statement and does not replace system-specific applicability analysis.

Policy area	ISO/IEC 42001:2023 management-system theme	EU AI Act theme
Context, scope and governance	Context of the organization; AI management system	Operator role, scope and risk-based regulatory framework
Leadership and policy	Leadership, AI policy, roles and authorities	Accountability and quality-management responsibilities for relevant systems

Policy area	ISO/IEC 42001:2023 management-system theme	EU AI Act theme
Risk and impact	Actions to address risks/opportunities; AI risk and system impact assessment	Risk management; data/fundamental-rights and high-risk requirements where applicable
Resources and competence	Resources, competence, awareness and communication	AI literacy and competent human oversight
Documented information	Creation, updating and control of documented information	Technical documentation, records, logs and traceability for applicable systems
Operational control	Operational planning and control; AI risk/impact assessment and treatment	Data governance, transparency, human oversight, accuracy, robustness and cybersecurity
Performance evaluation	Monitoring, measurement, analysis, internal audit and management review	Post-market monitoring and continuous evaluation for applicable systems
Improvement	Continual improvement; nonconformity and corrective action	Corrective action, serious incident handling and system updates for applicable systems

Reference Documents Used to Develop This Policy

- KOMGIS AI Security Policy - policy for protecting AI systems against prompt injection, data leaks, unauthorized access and unsafe automation.
- KOMGIS Data Privacy Policy - policy for handling client data in AI and software systems.
- KOMGIS AI Incident Response Plan - plan for handling wrong, dangerous, harmful or unauthorized AI outputs and actions.
- ISO/IEC 42001:2023 - Information technology - Artificial intelligence - Management system.
- Regulation (EU) 2024/1689 of the European Parliament and of the Council - Artificial Intelligence Act.

Policy interpretation questions shall be escalated to AI Governance / Technical Leadership and, where legal or regulatory interpretation is required, to qualified legal or compliance review.