# PageMind Security & Data Protection Overview

*Last updated: 2025-12-08*

*(Public overview – the binding terms are in the inAi/PageMind Terms, DPA and contracts.)*

> **IMPORTANT: NON-CONTRACTUAL INFORMATION**
> This document describes inAi's current security practices and data protection architecture for the PageMind service. It is provided for informational and transparency purposes only to assist current and prospective customers in their own risk assessments.
>
> **Binding Obligations:** The specific security measures, retention periods, and incident notification timelines that inAi is contractually required to maintain are **exclusively** set forth in:
>
> 1. The **Data Processing Agreement (DPA)** (specifically Annex 2); and
> 2. The **Master Service Agreement (MSA)** or applicable Order Form.
>
> In the event of any conflict between this Overview and the DPA or MSA, the DPA and MSA shall prevail. This document does not create independent rights or legal remedies. Any technical or organisational measures described in this Overview are indicative of our current implementation and risk management approach; they do not create additional contractual commitments beyond those expressly set out in the DPA, the Terms of Service, the MSA or applicable Order Form.

---

## 1. Purpose and scope

This document explains how inAi ("we") secures and processes customer data when you use PageMind, our catalog-ops automation platform for retailers, brands and marketplaces.

**Scope.** This overview describes the security and data-protection measures for the PageMind Retail catalog-ops workflows and the Verify.EU compliance module. It does not apply to other inAi products or services (such as Emplo), which have their own terms and notices.

It is designed to:

- Describe our **security and data-protection principles** for PageMind.
- Show, at a high level, **what data we process, where it lives and how long we keep it**.
- Clarify **roles and responsibilities** between you and inAi (shared responsibility model).([Microsoft Learn][1])

- Demonstrate alignment with **GDPR Article 32** (security of processing) and emerging EU AI regulation.([GDPR][2])

This is an **informational overview only**. It does **not** replace or override:

- Your **service agreement / MSA** with inAi.
- The **Data Processing Agreement (DPA)** that governs controller–processor obligations.
- Any **security annexes or SLAs** negotiated individually.

If there is any conflict, **the contract and DPA prevail**. Use of PageMind is always governed by the PageMind Terms of Service and the applicable DPA and MSA. This Overview is provided for transparency only and does not modify those terms or create any warranty or guarantee.

Nothing in this overview constitutes legal advice. It summarises our technical and organisational measures so that you and your advisers can perform your own assessment under the law and your internal policies.

---

## 2. Who we are and our regulatory baseline

### 2.1 Legal entity and jurisdiction

PageMind is provided by **INAI**, a French *SASU* (single-shareholder simplified joint-stock company) registered with the **RCS Lille Métropole** under number **987 977 386**, with registered office at **142 rue d'Iéna, 59000 Lille, France**.

Our corporate purpose is the **development and commercialization of software and SaaS services using artificial intelligence**, together with associated consulting, training and R&D.

### 2.2 EU and GDPR-first posture

- We are an EU-based company and design PageMind for **EU data-protection law from day one**.
- PageMind's core storage and orchestration are intended to run in an **EU-hosted virtual private cloud (VPC)**. However, specific AI inference tasks (e.g. LLM processing) may be routed to sub-processors located outside the EU (e.g. USA) where necessary for the service, always under appropriate GDPR transfer safeguards (such as Standard Contractual Clauses).
- Some processing steps may also be carried out by vetted **sub-processors** (including, where relevant, outside the EEA) under appropriate data-transfer safeguards, as described in our Data Processing Agreement (DPA) and published sub-processor list.

- We treat **GDPR Article 32** as a baseline and implement technical and organisational measures **designed to protect** the confidentiality, integrity, availability and resilience of processing systems, supported by encryption and regular testing.([GDPR][2])

### 2.3 Security & compliance foundations

We are not currently claiming any formal certification (e.g. ISO 27001) but we align our internal controls with recognised frameworks:

- **ISO 27001 Annex A control themes** (organisational, people, physical, technological) as a checklist for our ISMS design.([IT Governance][3])
- **GDPR security expectations** around encryption, access control, backup and regular testing.([GDPR Local][4])
- **EU AI Act risk-based approach.** Based on our current understanding of the EU AI Act, PageMind does not fall into any of the Annex III high-risk use cases. In typical retail catalog-ops deployments it would generally be treated as a limited-risk or minimal-risk system, but your own classification as a "deployer" will depend on how you use PageMind within your processes. We nevertheless apply AI-risk and logging practices inspired by the Act's requirements (transparency, logging, oversight).([Artificial Intelligence Act][5])

We are recognised by the **French State (DRIEETS Île-de-France)** as an **"economically innovative" project**, particularly for evidence-linked catalog workflows and EPREL-verified fields, and are incubated in EuraTechnologies' **Retail / e-commerce** vertical.

---

## 3. Data categories and roles

### 3.1 What kinds of data PageMind typically processes

PageMind is primarily a **B2B catalog-ops tool**. In normal use, it processes:

1. **Product and catalog data**
    - Supplier spec sheets and brochures (PDF, DOCX, images, Excel).
    - Product attributes (dimensions, power, materials, energy label data, etc.).
    - Titles, descriptions, bullets and marketing text.
    - Identifiers: SKUs, GTINs/EANs, internal product codes.
2. **Compliance-related information**
    - Energy-label and EPREL data for regulated products (e.g. appliances, TVs) via Verify.EU.
3. **Operational metadata**

    - File names, folder paths, run IDs.
    - Per-row trace (source file, page, processing steps, flags) and run reports with metrics, hashes and environment fingerprints.
4. **Account and usage data**
    - Business contact details for workspace users (name, email, role).
    - Configuration data (templates, glossaries, allowed lists, language settings).
5. **Limited log and diagnostic data**
    - System logs and error traces needed to maintain stability, security and audit trails.

PageMind is **not designed** to process end-consumer behavioural profiles, payment data or special categories of personal data. If you choose to upload such information, you remain responsible for ensuring a valid legal basis and for complying with the DPA and acceptable-use rules.

## 3.2 Personal vs non-personal data

Most data processed by PageMind in retail use is **business information about products**, not directly about identifiable individuals. However, the same technical and organisational safeguards apply when personal data is present (e.g. contacts in documents, internal staff names in logs).

We structure controls so that:

- **Product/catalog content** is treated as potentially sensitive commercial data.
- **Any personal data** in documents, accounts or logs is handled under GDPR standards for confidentiality, integrity, availability and resilience.([GDPR Local][4])

## 3.3 Controller / processor roles (high level)

Precise roles are defined in your **DPA**, but the typical model for PageMind Retail is:

- **You (the customer)** are usually the **data controller** for product/catalog data and any personal data you upload. You decide the purposes (e.g. building and publishing product pages) and means (templates, categories, languages).
- **inAi** acts as a **data processor** for that data, providing hosted processing under your instructions (configuration, templates, workflows) and only using it to:
    - Run and improve PageMind for your workspace;
    - Provide support and help maintain security;
    - Comply with legal obligations.

For some narrow categories (e.g. **billing metadata, aggregated usage metrics, security logs**), inAi may act as an **independent controller**; this is clarified in the DPA and Terms.

## 4. High-level data flow in PageMind

This section describes how a typical PageMind run handles your data from ingestion to deletion.

### 4.1 Ingestion

1. **You choose the inputs**
    - You place supplier files (PDF, DOCX, Excel, images, etc.) into an input folder or workspace and configure a run: target template, languages, glossary and allowed lists.
    - You remain responsible for the legality and appropriateness of the content you upload (e.g. ensuring you have rights to use it and it does not include prohibited data).
2. **Secure transfer**
    - Files are transmitted over **encrypted channels (TLS)** between your browser/agent and our EU VPC.
3. **Temporary staging**
    - Files are stored in **workspace-scoped storage** in our EU VPC. Access is restricted to that workspace and to a small number of security-cleared staff under strict need-to-know rules.

### 4.2 Processing

PageMind runs a deterministic pipeline from raw files to catalog outputs and reports:

1. **Parsing and OCR**
    - Digital documents are parsed directly; scanned PDFs and images go through OCR.
    - OCR errors are converted into structured retry reasons rather than silent failures where possible.
2. **Product detection & structuring**
    - The engine identifies product “chunks” and builds an internal product record per SKU.
    - Heuristics are used for complex catalogs; gaps here are tracked and improved as part of our roadmap.
3. **Attribute extraction & normalization**
    - Attributes (e.g. color, size) are extracted and normalized against **allowed lists**, dropping values when evidence is missing or ambiguous.

- Format anomalies and low-confidence extractions are flagged for QA.

4. **Text generation & translation**
    - Titles and descriptions are generated in the configured languages under glossary control.
    - We instruct models **not to invent facts** and to stay within the information present in your documents, but final responsibility for reviewing outputs before publication remains with you.
5. **Verify.EU (compliance autopilot, when enabled)**
    - For regulated energy-label products, Verify.EU compares supplier data with official EPREL records and applies a strict **"no evidence, no publish"** rule for compliance fields, creating an evidence pack per SKU for audit.
6. **Output & QA**
    - The pipeline emits:
        - Main template-aligned CSV;
        - QA review CSV (rows needing human attention);
        - Retry list with reason codes;
        - Run report with metrics, hashes, config signatures and environment snapshot.

## 4.3 Storage, retention and deletion

We design retention to balance auditability with minimisation:

- **Active workspace**
    - Input files, outputs and run reports are stored in workspace-scoped storage for as long as your workspace is active, subject to your contract and storage limits.
- **Backups and logs**
    - Backups and system logs are kept for limited periods to support resilience and security. As of the date of this Overview, our target retention window for such backups and logs is approximately thirty (30) days after workspace deletion, unless a longer period is required by law or explicitly agreed in the DPA or an Order Form. These indicative timeframes may evolve as our infrastructure and risk assessments change.
- **Deletion**
    - When you delete a workspace or dataset (or request deletion), we schedule deletion from active storage and make sure associated backups expire within the defined window.

- We may retain minimal records necessary for billing, legal compliance and abuse prevention, as permitted under GDPR.([ICO][6])

### 4.4 Use of data for training and improvement

- We **do not** use your catalog content to train public, general-purpose models that are made available to unrelated customers.
- We may use **aggregated and, where possible, pseudonymised telemetry and error patterns** to maintain and improve PageMind's reliability and security (for example, better detection of mis-splits or OCR failures), in line with the purposes and safeguards described in the DPA.
- Where we rely on third-party AI APIs, we select providers that offer contractual and technical options consistent with our data-protection obligations and configure services, where possible, so that your prompts and data are **not used to train their general models**. The current list of such sub-processors and their locations is published in our sub-processor list referenced in the DPA. To the maximum extent permitted by applicable law, we are not responsible for any processing that such third-party AI providers carry out for their own independent purposes as controllers, outside our documented use of their services on your behalf.

---

## 5. Technical and organisational security measures

We design PageMind's security controls around the principles of **defence in depth**, **least privilege** and **auditability**, while remaining realistic about risk and avoiding absolute claims.

### 5.1 Security governance

- **Responsibility and oversight**
    - Overall responsibility for information security sits with the **President/CEO of INAI**, supported by technical leads responsible for infrastructure and application security.
    - Security is treated as a core part of product design for PageMind, not a separate afterthought.
- **Policies and processes**
    - inAi is **establishing and maintaining** internal security policies covering areas such as access control, acceptable use, change management, incident response and vendor management. At our current size, these policies are intentionally concise and may evolve as the company and PageMind grow.
    - We review these policies from time to time and adjust them as laws, standards and our infrastructure evolve.

- **Risk management**
    - We carry out **risk-based assessments** of threats to the confidentiality, integrity and availability of PageMind and adjust controls as needed. Particular attention is paid to:
        - exposure of catalog and supplier data;
        - unauthorised access to evidence packs and compliance artefacts;
        - abuse of the platform to process prohibited data/categories;
        - dependency on third-party AI and cloud providers.

Nothing in this document should be interpreted as creating an assurance of absolute security; instead, it describes how we organise ourselves to manage risk in a proportionate way.

## 5.2 Infrastructure and network security

- **EU-hosted VPC with restricted connectivity**
    - PageMind's core infrastructure is deployed in an **EU-hosted virtual private cloud**. Network security groups and firewalls restrict inbound and outbound connections to what is necessary to operate the service and to communicate with approved sub-processors.
- **Network segmentation and filtering**
    - Internal services are segmented so that only components that must communicate with each other can do so.
    - Firewalls and security groups restrict inbound and outbound connections, with default-deny policies for non-essential traffic.
- **Hardened infrastructure**
    - Systems follow baseline hardening practices (minimal services, regular patching, restricted admin access).
    - Administrative access to production infrastructure is limited to a small number of authorised individuals and is subject to strong authentication and logging.
- **Environment isolation**
    - Development, testing and production workloads are separated; test environments do not use production customer data unless explicitly agreed and safeguarded.

## 5.3 Identity and access management

- **Authentication**

    - Access to production infrastructure and administrative tooling at inAi requires strong authentication mechanisms; multi-factor authentication (MFA) is mandatory for internal administrative and production-access accounts.
    - In the current private-beta deployments of PageMind, customer access is typically controlled through account-level credentials or your own identity provider/SSO where integrated. We strongly recommend that you enforce MFA and appropriate device and network security for those accounts on your side.
- **Authorisation and least privilege**
    - During the current beta, PageMind workspaces are operated in a **single-tenant model per customer**. Within a given workspace, access is limited to the accounts and integrations that you or we explicitly configure; there is **no fine-grained in-product role-based access control (RBAC) yet**.
    - Internal staff access to customer data is restricted to a minimal, security-cleared subset of personnel on a need-to-know, time-limited basis and is always logged. Access for development and support is provided only where necessary and is revoked when no longer needed.
- **Account and access lifecycle**
    - You are responsible for deciding who within your organisation is authorised to access PageMind and for managing the lifecycle of those accounts and any connected systems (for example, SSO identities, VPN accounts or desktop access).
    - When we later introduce finer-grained RBAC features within PageMind, this overview and the DPA will be updated to describe how those controls work and how responsibilities are shared.

## 5.4 Data protection measures

- **Encryption in transit and at rest**
    - All communications between customer clients and PageMind use **encrypted channels (TLS)**.
    - Customer content, configuration data and run metadata stored in our infrastructure are encrypted at rest using industry-standard encryption mechanisms.
- **Data minimisation and separation**
    - PageMind only stores the data required to execute workflows, provide outputs, maintain logs and meet legal obligations.
    - Where technically feasible, diagnostics and telemetry are pseudonymised or aggregated to limit exposure of personal data.

    - Customer workspaces are logically separated; data from one workspace is not visible to users of another.
- **Backups and retention**
    - We maintain backups for resilience and disaster recovery; backup retention is limited, and backups are purged within approximately **30 days** after workspace deletion, except where a longer period is required by law or contract.
    - Customers are responsible for storing any exports (CSV files, evidence packs, reports) that they download into their own systems.

## 5.5 Application security and secure development

- **Secure development practices**
    - We aim to follow a secure software development lifecycle appropriate to our size, which includes, where practicable:
        - peer review of significant changes;
        - automated dependency and vulnerability scanning on key components;
        - staged rollouts and controlled deployments for production changes;
        - regression and integration testing for core PageMind pipelines.
- **Protection against common web threats**
    - PageMind is designed with defences against common classes of vulnerabilities (such as injection, broken authentication, insecure direct object references and cross-site request forgery) guided by recognised best-practice lists such as OWASP.
- **Configuration management**
    - Critical runtime configurations (such as model routing, glossary rules, allowed lists and template definitions) are versioned and can be traced through run reports, enabling safe evolution and rollback when needed.

We may adjust specific technical implementations over time (for example, changing cloud providers or libraries) while maintaining a comparable or higher level of protection.

## 5.6 Logging, monitoring and audit trails

- **Operational logging**
    - PageMind records operational logs for:
        - workflow execution (runs, stages, success/failure);
        - system health and performance;
        - error and exception events;

        - security-relevant events (authentication attempts, privilege changes).
- **Run-level audit trails**
    - Each catalog run produces a **run report** that includes:
        - counts and fill-rates;
        - flags and anomaly summaries;
        - hashes of inputs and outputs;
        - configuration signatures (prompt/model versions, glossary and allowed-list hashes);
        - environment fingerprints for reproducibility.
- **Access logging**
    - Access to customer workspaces, runs and evidence packs is logged for audit and forensic purposes.
- **Monitoring and alerting**
    - Automated and manual monitoring are used to detect unusual patterns (for example, repeated failures or abnormal traffic) and to trigger alerts for investigation.

These logs and reports are designed to support both **security monitoring** and **business/compliance audits** while respecting data minimisation and retention policies.

---

## 6. AI-specific safeguards and auditability

PageMind relies on AI components (including LLMs and OCR engines) to analyse supplier documents and generate or transform text. Our design goal is to **constrain and evidence** AI behaviour rather than rely on unstructured "best-effort" generation.

### 6.1 Evidence-first design

- **"No evidence, no publish" in compliance modules**
    - In the Verify.EU module, every compliance-critical field (such as energy class and consumption metrics) must be backed by an explicit evidence object referencing a trusted source (EPREL or supplier documentation).
    - Where sufficient evidence is not available or is contradictory, the field is left blank and surfaced as a QA or retry item.
    - This hard rule currently applies **only within the Verify.EU module**. In the broader catalog pipelines, PageMind is designed to surface missing or weak evidence and to avoid inventing facts, but you are still expected to review outputs and decide what to publish.
- **Per-attribute evidence and traceability**

    - For key attributes such as color and size, PageMind can attach short quotes and page ranges from the originating files as evidence.
    - The main CSV includes at least per-row trace columns such as `source_file`, `source_page` and `steps_used`, enabling teams to reconstruct where values came from.
    - During the current beta, some rows may still be emitted without complete trace information; in those cases, run-level reports and flag columns are intended to highlight missing trace so that you can focus manual review where it is most needed.
- **Separation of extraction and generation**
    - Where possible, PageMind separates **fact extraction** from **text generation**:
        - extraction pipelines focus on structured attributes and numeric values,
        - generation pipelines focus on human-readable titles and descriptions based on extracted facts.

## 6.2 Determinism and reproducibility

- **Deterministic pipelines**
    - PageMind's catalog runs follow a fixed, auditable pipeline rather than arbitrary prompts. Every run is structured around the same sequence of steps (ingestion → extraction → translation → QA → export).
- **Run fingerprints**
    - Each run report contains hashes of inputs and outputs, along with configuration and environment fingerprints, so that results can be reproduced or differences between runs can be explained.
- **Controlled variability**
    - Some generative steps (e.g. phrasing of descriptions) may involve limited model randomness for naturalness.
    - For compliance-critical and numeric fields, we aim to minimise variability across runs by applying stricter rules, structured extraction and evidence checks.

No system using probabilistic models will produce identical outputs in every scenario, but these mechanisms are intended to keep behaviour **predictable, explainable and auditable**.

## 6.3 Model usage and training policy

- **Processing scope**
    - AI models (including third-party LLMs) are used to:

        - parse and interpret documents;
        - extract and normalise attributes;
        - generate and translate catalog text under glossary control;
        - support QA and anomaly detection.
- **Use of customer content for training**
    - Customer content is processed to provide PageMind's services and to maintain security, quality and reliability.
    - Unless explicitly agreed in writing, we **do not use customer content to train foundation models that are made available to other customers as generic AI services**.
    - Aggregated statistics and error patterns derived from usage (for example, language mix, typical failure modes) may be used, in **aggregated and where possible pseudonymised form**, to improve PageMind's algorithms and safeguards without exposing identifiable content outside the scope described in the DPA.
- **Third-party AI providers**
    - Where we use third-party AI APIs, we select providers that offer contractual and technical controls aligned with our data-protection obligations and configure services, where possible, **not to use customer prompts or data to train their general models**.
    - Details of such providers and their locations are made available in our sub-processor list and DPA.

## 6.4 Human oversight and responsibility

- **Tool, not decision-maker**
    - PageMind is an automation tool that **assists** your teams; it does not replace your legal, compliance or product-content functions.
    - The final decisions to approve content, publish catalog entries or rely on Verify.EU for compliance remain entirely with you.
- **Human-in-the-loop workflows**
    - The system surfaces:
        - a QA review CSV for rows requiring human attention;
        - retry lists with reason codes;
        - flags for missing, ambiguous or anomalous data.
    - You are expected to review these outputs and implement internal controls proportionate to the risk class of products and markets involved.
- **AI transparency**

- PageMind is designed to make it clear when AI has been used and what it has done: runs, flags, evidence links and configuration versions are visible at the level required for internal audits and external enquiries.

### 6.5 Limits and disclaimers

- While we implement safeguards to reduce hallucinations, mis-extractions and inconsistent outputs, **no AI system is error-free**.
- PageMind does **not** provide legal advice and cannot promise that catalog content or compliance information produced using the service will always be complete, accurate or compliant with all applicable laws or platform rules.
- You must maintain appropriate human review and validation procedures and remain responsible for the accuracy and legality of the content you publish, including any content created or assisted by PageMind.

## 7. Privacy and data subject rights (high level)

This section summarises how PageMind supports GDPR obligations. Full details appear in the **Privacy Policy** and **Data Processing Agreement (DPA)**.

### 7.1 Lawfulness, fairness and transparency

- **Controller responsibility**
  - As controller, you are responsible for ensuring a **lawful basis** for any personal data embedded in documents or configurations you upload to PageMind (e.g. legitimate interest in processing supplier contacts or internal staff names).
  - You remain responsible for informing the relevant data subjects where required (e.g. employees, suppliers).
- **Processor role of inAi**
  - When acting as processor, inAi processes personal data **only on your documented instructions** as set out in the DPA and configuration you choose in PageMind.
  - We do not decide new purposes for your catalog/supplier data; we use it to:
    - execute your catalog workflows;
    - secure and maintain PageMind;
    - meet legal obligations (e.g. fraud or abuse investigations, accounting).
- **Transparency and documentation**
  - We are **building and maintaining** documentation of PageMind's processing activities, categories of data and sub-processors and will expand that documentation as the service evolves. Where required, we can make

relevant parts of this documentation available under NDA to support your own assessments.
    - Where we act as controller (for example, for account management, billing or website analytics), we explain those activities in our Privacy Policy.

## 7.2 Supporting data subject rights

Under GDPR, data subjects have rights of access, rectification, erasure, restriction, portability and objection in relation to their personal data.

- **Your role as controller**
    - You are responsible for handling and responding to data-subject requests concerning data for which you are the controller (including catalog content and documents you upload).
    - You determine whether a request is valid, what data it concerns, and how to respond.
- **Our support as processor**
    - inAi will assist you, as required by GDPR and the DPA and **taking into account the nature of the processing and our technical capabilities**, in:
        - locating relevant personal data within PageMind;
        - implementing deletion or restriction in PageMind where technically feasible;
        - providing information about systems and logs where needed to document your response.
          This assistance may be limited where data is stored only in backups or log archives that are not reasonably searchable without disproportionate effort; the DPA explains how those situations are handled.
    - Where a request is addressed directly to inAi for data we process on your behalf, we will, where appropriate, **forward the request to you** and follow your instructions.
- **Response times and cooperation**
    - The DPA sets out the timeframes and conditions under which we provide assistance (e.g. reasonable notice and scope), ensuring that you can meet your own legal deadlines (such as GDPR's one-month response target).

## 7.3 Special categories, children's data and prohibited content

PageMind is designed for **B2B product and catalog data**, not for special categories of personal data or children's data.

- **Special categories and criminal data**
    - You must **not** knowingly upload special categories of personal data (for example, health, biometric, political or religious data) or criminal-offence data into PageMind, unless we have entered into a **specific written agreement** providing additional safeguards for that processing.
    - If you nonetheless upload such data without a specific agreement, you remain responsible for ensuring a lawful basis and appropriate additional safeguards under applicable law. We will treat any such data using the same technical and organisational security measures described in this overview, but we may delete, restrict or suspend processing of it and are not obliged to implement sector-specific or higher-risk safeguards unless expressly agreed in writing.
- **Children's data**
    - PageMind is not intended to process children's personal data. You should not use PageMind for activities that involve children's data without a tailored contractual arrangement and additional safeguards.
    - Where we become aware that a workspace is being used in clear breach of this rule, we may suspend or restrict access to that workspace and request that you remove the offending data, in line with the Terms and Acceptable Use Policy.
- **Illegal and harmful content**
    - You must not upload content that is unlawful, infringing, abusive or malicious (such as malware or deliberately harmful prompts). Such use may lead to suspension or termination under the Terms and Acceptable Use Policy.

## 7.4 Cookies and tracking (website and app)

- **Website and app telemetry**
    - Our marketing website and, where applicable, PageMind interfaces may use cookies and telemetry to understand usage, secure sessions and improve the service.
    - Non-essential cookies (such as analytics or marketing cookies) are only used with your consent where required by law, via our cookie banner and settings.
- **Detailed information**
    - Details of cookie types, purposes and retention periods are set out in our **Cookie Policy** and **Privacy Policy**.

---

## 8. Customer responsibilities and shared responsibility model

Security and compliance in PageMind rely on a **shared responsibility model**. This document describes what we do on our side; this section clarifies what you must do on yours.

### 8.1 Data responsibility and lawful use

You remain primarily responsible for the **content and legality of the data** you have PageMind process. This includes:

- ensuring you have all necessary rights, licences and permissions for supplier documents and catalog data;
- ensuring a valid legal basis under GDPR where personal data is present;
- not using PageMind to process prohibited or inappropriate data categories (as described in Section 7.3);
- ensuring data quality (e.g. removing obviously incorrect or irrelevant records before ingestion).

PageMind will not and cannot verify the legal basis on which you collected or use your underlying data; that assessment remains with you as controller.

### 8.2 Output verification and regulatory compliance

PageMind aims to **reduce manual work and surface issues**, but it does not replace your internal review and compliance functions. You remain solely responsible for:

- **Accuracy and completeness of published content**
    - Verifying titles, descriptions, attributes and images before they are published in your PIM, e-commerce platform or other channels.
    - Using the QA CSV, retry lists and flags provided by PageMind to focus human review on the riskiest or most ambiguous items.
- **Regulatory and platform compliance**
    - Ensuring that product information, labels and disclosures meet all applicable laws and platform policies in the markets where you operate (e.g. EU energy labelling rules, national consumer law, marketplace rules).
    - Using Verify.EU as a **supporting tool** only, not as a definitive legal source, and reconciling discrepancies between EPREL data, supplier docs and your own internal records.
    - Complying with any platform-specific rules or registry terms that apply to services used via PageMind (for example, EPREL conditions of use or marketplace terms of service); PageMind and Verify.EU do not assume those obligations for you.

- **Risk-appropriate controls**
    - Designing and enforcing internal approval flows (e.g. which teams sign off catalog content before publish) proportionate to the risks of your products and jurisdictions.

No matter how PageMind is configured, **you are the party that ultimately decides what is shown to consumers and regulators**; you carry legal responsibility for those decisions and their consequences.

To the maximum extent permitted by applicable law, we do not accept responsibility for any non-compliance, fines, penalties or third-party claims resulting from catalog content that you create, approve, publish or rely upon, even where such content was assisted by PageMind.

## 8.3 Security on your side

Our controls protect PageMind's infrastructure and data within it, but you control what happens at the edges:

- **Account and access management**
    - Create, manage and revoke user accounts and roles in PageMind so that only appropriate staff can access particular workspaces and runs.
    - Enforce strong authentication (including MFA where available) and appropriate security on your own identity provider or SSO.
- **Endpoint and integration security**
    - Secure the devices and networks from which your users access PageMind (patching, endpoint security, safe browsing practices).
    - Secure any integrations between PageMind and your systems (e.g. PIM, storage, internal APIs), including API keys, connection secrets and access policies.
- **Handling of exports**
    - Once you download CSVs, evidence packs or reports from PageMind, they become part of your own systems and are governed by your own security and retention policies.
    - You are responsible for preventing unauthorised access, loss or disclosure of these exported files.
    - Without limiting the contractual limitations of liability, inAi is **not responsible** for any loss, corruption or unauthorised access occurring in your systems or in third-party platforms after data has been exported from PageMind.

If your environment is compromised (for example, a stolen laptop or compromised SSO account), an attacker may gain access to PageMind via your credentials; this is outside of inAi's control and responsibility.

## 8.4 Platform terms, acceptable use and anti-abuse

Use of PageMind is subject to the **PageMind Terms of Service**, the **DPA** and the **Acceptable Use Policy**. In particular, you must not:

- attempt to bypass security or technical restrictions, such as rate limits or access controls;
- probe or test PageMind for vulnerabilities except under an agreed responsible-disclosure or testing programme;
- use PageMind to process illegal content, infringing material or data you are not authorised to handle;
- abuse EPREL or other third-party services accessed via PageMind in ways that violate their terms.

We may monitor for signs of abuse or misuse and, where necessary, **suspend or restrict access** to protect the security and integrity of PageMind or to comply with legal obligations.

## 8.5 Allocation of risk at a high level

Detailed allocation of risk, warranty and liability is set out in the PageMind Terms and any MSA you sign. At a high level:

- inAi is responsible for:
    - providing PageMind with reasonable care and skill;
    - implementing the technical and organisational security measures described in this overview (as they evolve);
    - complying with its obligations as a processor and, where applicable, as an independent controller.
- You are responsible for:
    - your choice to use PageMind, your configuration of its workflows, and the data you send to and publish from PageMind;
    - compliance with laws, regulations and third-party terms that apply to your business and products;
    - implementation of appropriate human review, approval and quality-control processes.

This shared responsibility model is essential to keeping risk manageable for both parties.

For clarity, in any case of discrepancy between this Overview and the allocation of responsibilities, warranties and liability in the PageMind Terms of Service or a signed MSA, the latter shall prevail.

## 9. Incident management and breach notification

We have established, and intend to maintain, a formal Incident Response Plan (IRP) inspired by NIST and ENISA guidelines, and we review and adjust this plan from time to time as our infrastructure and risk profile evolve. While specific notification timelines are governed by your DPA, our operational approach follows these phases:

### 9.1 Detection and Classification

- **Sources:** Incidents are detected via automated monitoring (infrastructure metrics, error rates), internal staff reports, and notifications from sub-processors.
- **Triage:** Upon detection, events are triaged by the Security Lead to determine if they constitute a "Security Incident" (compromising confidentiality, integrity, or availability) or a "Personal Data Breach" under GDPR.
- **Severity:** Incidents are classified by severity (Critical, High, Medium, Low) based on the scope of data affected and impact on service continuity.

### 9.2 Response Lifecycle

1. **Containment:** Immediate steps to limit impact (e.g., revoking compromised credentials, isolating network segments, or temporarily suspending a workspace).
2. **Investigation:** Forensic analysis of logs and run reports to determine root cause and the extent of any data exposure.
3. **Eradication & Remediation:** Removing the root cause (e.g., patching a vulnerability, removing malicious files) and verifying system integrity.
4. **Recovery:** Restoring services and data from secure backups.
5. **Post-Incident Review:** A mandatory retrospective analysis for all High/Critical incidents to update policies and prevent recurrence.

### 9.3 Notification Commitments

- **To Customers (Controller Notification):** If a Personal Data Breach affects your data, we will notify you **without undue delay** after becoming aware of the breach, in accordance with the DPA. We do not wait for a full forensic conclusion to provide an initial alert if risk is confirmed.
- **To Authorities/Subjects:** As a processor, we assist you in meeting your obligations to notify supervisory authorities (e.g., CNIL) or data subjects. We generally do not notify them directly unless required by law or agreed in the DPA.
- **Transparency:** Notices will include the nature of the breach, likely consequences, and measures taken to address it.

## 10. Business continuity and disaster recovery

PageMind is designed to support the continuity of your catalog operations, but no online service can promise zero downtime or data loss. This section explains our general approach; specific uptime commitments and remedies, if any, are defined in the applicable SLA or Terms.

### 10.1 Availability objectives

Unless a specific Recovery Time Objective (RTO) is defined in an Order Form or SLA, the following describes general targets and design principles rather than binding service levels.

- **Service design**
    - Core components of PageMind are deployed in an EU cloud environment designed for high availability, with redundancy at the infrastructure level (multiple availability zones where supported).
    - Stateless services are designed to be horizontally scalable; stateful components (e.g. databases, storage) use provider-level redundancy features where available.
- **Targets and expectations**
    - We aim for high availability suitable for production catalog operations; precise numerical availability targets and any corresponding service credits are defined in your contract or SLA where applicable.
    - Maintenance windows and major changes are, where possible, planned to minimise customer impact.

### 10.2 Backups, replication and data recovery

- **Backups**
    - We perform regular backups of critical data stores so that we can restore data in the event of corruption or catastrophic failure.
    - Backups are encrypted and stored within the EU, subject to the same access controls as primary data.
- **Replication**
    - Where supported, we use storage and database replication mechanisms to improve resilience against hardware or zone failures.
- **Recovery procedures**
    - We maintain documented procedures for restoring services and data from backups following incidents such as infrastructure failures, data corruption or accidental deletion in our environment.

    - Recovery from backups is prioritised by impact and feasibility; in some scenarios, partial restoration (e.g. specific workspaces or time windows) may be more appropriate than full-environment rollback.
- **Limitations**
    - Backups and replication are not a substitute for your own resilience strategies; once data leaves PageMind (for example, via CSV export into your PIM or other systems), you must put in place your own backup and recovery mechanisms for that downstream data.

## 10.3 Continuity planning and testing

- **Continuity planning**
    - We are **documenting and iteratively improving** internal business-continuity and disaster-recovery (BC/DR) procedures that identify critical PageMind components, key dependencies and recovery priorities.
    - These procedures focus on realistic scenarios for a company of our size, such as cloud-provider incidents, regional disruptions and loss of key internal systems.
- **Testing and improvement**
    - As PageMind moves beyond private beta, we plan to exercise selected continuity and recovery procedures through simulations or controlled failover tests and to feed the lessons learned into both our technical setup and our documentation.

## 10.4 Customer responsibilities for continuity

Your continuity posture depends not only on PageMind but also on how you integrate and rely on it. You are responsible for:

- assessing the criticality of PageMind within your own operations and designing appropriate fallback plans (for example, temporary manual workflows or alternative data sources during outages);
- implementing monitoring and alerts on your side for failures in downstream processes (e.g. imports into your PIM or e-commerce platform);
- ensuring that you keep copies of critical exports and configurations needed to rebuild or repopulate systems if necessary;
- factoring PageMind's documented availability characteristics and contractual SLAs into your broader business-continuity and disaster-recovery planning.

inAi cannot be responsible for the consequences of outages or failures in systems outside our control (such as your PIM, websites, networks or third-party platforms), even where they interact with PageMind.

# 11. Compliance posture and roadmap

## 11.1 Current posture

PageMind's security and data-protection posture is built around:

- **EU and GDPR-first design**
    - Core processing in an **EU-hosted VPC** with defined deletion windows and strong access controls, plus vetted sub-processors under appropriate data-transfer safeguards as set out in the DPA.
    - Architecture and processes **designed with GDPR Article 32 in mind** (appropriate technical and organisational measures, including encryption, resilience and periodic testing proportionate to our size and risk profile).([Digital Strategy][1])
- **Evidence-centric catalog and compliance workflows**
    - "No evidence, no publish" rule for Verify.EU compliance fields.
    - Per-row trace (source file, pages, steps) and run-level reports with hashes and configuration fingerprints for reproducibility and audit.
- **B2B focus and minimal personal data**
    - Primary focus on product and catalog data for retailers and brands, with personal data typically limited to account and incidental contact information.
- **Innovation recognised at national and ecosystem level**
    - INAI is recognised by the French State (DRIEETS Île-de-France) as an economically innovative project, with PageMind's evidence-linked workflows and EPREL use at its core, and is incubated at EuraTechnologies in the Retail vertical.

This overview describes a **living security and compliance posture**. Controls evolve over time as laws, standards, threats and PageMind itself evolve.

## 11.2 Standards and frameworks we align with

We do not claim formal certification unless explicitly stated in separate documentation. Instead, we **align** our controls with relevant European regulations and commonly used frameworks:

- **GDPR / EU data-protection law**
    - Article 32 (security of processing), Article 28 (processor obligations) and associated guidance on encryption, access control, logging, and breach notification.([Digital Strategy][1])

- **EU AI Act – risk-based approach**
    - Based on our current reading of the EU AI Act, PageMind does not fall into any of the Annex III high-risk categories. In typical retail catalog use it would generally be treated as a limited-risk or minimal-risk AI system, subject mainly to transparency-style obligations. This assessment is **informational only** and does not replace your own classification as a deployer in your specific context.([Artificial Intelligence Act][2])
    - Nevertheless, we adopt practices inspired by high-risk obligations where appropriate (logging, documentation, human oversight, transparency on AI use), and we intend to document our own internal risk assessment for PageMind as the Act and related guidance are finalised.([EU AI Act][3])
- **NIS2-style cybersecurity risk management**
    - Although PageMind may not itself be an "essential" or "important" entity as defined by NIS2, we take guidance from its emphasis on:
        - risk management and governance;
        - incident handling and reporting;
        - business continuity and disaster recovery;
        - supply-chain and sub-processor security.([Digital Strategy][4])
- **Industry frameworks and best practices**
    - We use ISO 27001 control themes and ENISA technical guidance as reference points for designing our policies and controls, particularly for access control, logging, change management and incident handling.([Hyperproof][5])

These references are **orientation points**, not marketing claims of full compliance or certification. The binding legal obligations and allocations of responsibility are defined in your contract and the DPA.

## 11.3 Roadmap and regulatory evolution

The regulatory environment for AI and cybersecurity is evolving rapidly. Our roadmap includes:

- **AI Act phase-in**
    - The EU AI Act is entering into force progressively, with transparency obligations for certain AI systems (including "limited-risk" systems) and specific rules for general-purpose AI models.([Digital Strategy][1])
    - We monitor the AI Act's implementing measures, guidance and emerging **Codes of Practice** that aim to help providers and users of AI comply in a risk-based, pragmatic way.([CMS Law][6])

- **NIS2 transposition and guidance**
    - NIS2 significantly raises expectations on cybersecurity risk management, incident reporting and supply-chain security across many sectors.([Digital Strategy][4])
    - We track national transposition and sectoral guidance (e.g. from ENISA and national CSIRTs) to help keep PageMind's controls consistent with the direction of travel even where NIS2 may not apply directly to us.([ENISA][7])
- **Internal governance and documentation**
    - Periodic reviews of our data-protection impact assessments (DPIAs) and risk assessments for PageMind.
    - Iterative enhancement of logging, evidence and documentation to support customers' own AI and cybersecurity governance obligations, including internal committees, auditors and regulators.

Our commitment is not to "freeze" PageMind's security at a single point in time, but to **maintain and improve** it in line with realistic, risk-based expectations and evolving EU regulation.

---

## 12. Versioning, contact and use of this overview

### 12.1 Versioning and changes

- **Version and date**
    - This document is issued with a **version number and date** shown near the top.
    - When we make material changes to our security or data-protection posture, we update this overview and adjust the version and date accordingly.
- **Change log**
    - For significant changes (for example, new classes of sub-processors, major architecture changes, or materially different controls), we maintain a brief change log describing what has changed at a high level.
    - Where required by the DPA or contract, we will notify customers of material changes (for example, to the sub-processor list or to measures that materially affect the level of protection).
- **Relationship to contracts**
    - This overview is **non-contractual**. In case of discrepancy or conflict between this document and the Terms of Service, MSA, DPA or other binding contract, the contractual documents prevail.

- The applicable law and jurisdiction for any dispute are those specified in your Terms of Service or MSA; this overview does not set or modify them.

## 12.2 Contact and further information

- **Security and privacy enquiries**
  - For security or privacy questions about PageMind, customers can contact us using the contact details published in our Legal Notice or on the contact page of our website.
  - Depending on the nature of the request, we may ask you to use a designated security or privacy contact channel (for example, a dedicated mailbox) so that we can triage and respond efficiently.
- **Additional documentation**
  - Under NDA or as part of procurement and security-review processes, we may provide additional information such as:
    - a copy of our standard DPA and sub-processor list;
    - responses to security or privacy questionnaires;
    - more detailed architectural or operational descriptions relevant to your risk assessment.
- **Responsible vulnerability disclosure**
  - If you believe you have discovered a security vulnerability affecting PageMind, we ask that you report it to us privately using the contact details on our website, providing enough information to reproduce the issue.
  - We request that you refrain from public disclosure until we have had a reasonable opportunity to investigate and remediate.

## 12.3 How to use this overview

This overview is intended to:

- help you understand, at a high level, **how PageMind approaches security and data protection**;
- support your internal **risk, security and compliance reviews**;
- complement, but not replace, the more detailed and binding provisions in the **Terms of Service, MSA and DPA**.

You should read this document together with:

- the **PageMind Terms of Service** or your negotiated MSA;
- the **Data Processing Agreement** governing GDPR roles and obligations;
- our **Privacy Policy** and **Cookie Policy**;
- the **Acceptable Use Policy** and the **sub-processor list**.

Taken together, these documents provide the full picture of how PageMind operates, the measures we implement, and how responsibilities are allocated between you and inAi.