



# SettlDay

SETTLEMENT READINESS PLATFORM

GET IT DONE

---

## AI Governance Pack

Version dc151e3  
Generated 2026-04-14

SettlDay – A joint initiative by Kommalpha AG and amalics GmbH

[settl.day](https://settl.day)

# Contents

## 1. AI Governance

- 1.1 AI Features
- 1.2 EU AI Act Classification
- 1.3 Transparency
- 1.4 Data Minimisation Pipeline
- 1.5 AI Provider Architecture
- 1.6 MiFID II Considerations
- 1.7 Record-Keeping and Auditability
- 1.8 Local Embeddings
- 1.9 Data Sent to the AI Provider
- 1.10 Art. 22 – Automated Decision-Making

## 2. Data Flow Architecture

- 2.1 Quick Facts
- 2.2 Data Flow Steps
- 2.3 Compliance Framework
- 2.4 Data Flow Summary
- 2.5 EU Data Residency

# 1. AI Governance

<b>Document ID</b>	TC-AI-001
<b>Classification</b>	Public – Trust Center

## 1.1 AI Features

SettlDay uses artificial intelligence in two areas:

Feature	Description
<b>Paid report generation</b>	Multi-stage generation of personalised T+1 settlement readiness reports, including executive summary, category analysis, remediation roadmap, cost analysis, regulatory mapping, and peer benchmarking
<b>RAG research assistant</b>	Question-answering interface over a curated T+1 regulatory research library, with source attribution for every response

Both features are powered by **AWS Bedrock** (EU regions only). Neither feature makes autonomous decisions with legal effect. All AI outputs are advisory and presented for human review.

## 1.2 EU AI Act Classification

Criterion	Assessment
<b>Autonomous decision-making</b>	No – all AI outputs are advisory; assessment scoring is rule-based, not AI-driven
<b>High-risk category (Annex III)</b>	Not directly listed; financial sector AI for compliance advice is a grey area
<b>Classification</b>	<b>Limited risk</b> – transparency obligations apply
<b>Rationale</b>	AI generates reports for human review; no credit scoring, trading, or automated compliance determinations

SettlDay is actively monitoring EU AI Act developments and will seek formal legal opinion on classification as regulatory guidance matures.

## 1.3 Transparency

SettlDay meets the transparency obligations under EU AI Act Art. 50:

Requirement	Implementation
<b>AI-generated content disclosure</b>	All AI-generated content is explicitly marked as AI-generated; an "AI-Generated Content Disclosure" box appears on the first content page of every report
<b>AI disclaimers</b>	Every report includes a disclaimer explicitly referencing AI generation on the back page
<b>Model identification</b>	All model identity is returned in research assistant responses and logged per report generation stage
<b>Advisory framing</b>	All AI outputs are explicitly framed as advisory – not as investment advice, binding compliance determinations, or autonomous decisions

## 1.4 Data Minimisation Pipeline

SettlDay's core differentiator is a multi-layer data minimisation pipeline that ensures personal and confidential data never reaches the AI provider. All transformations occur entirely on SettlDay's own EU infrastructure before any data is transmitted.

### 1.4.1 Layer 1 – PII Detection

Personal identifiers in free-text fields and document excerpts – including names, email addresses, and other recognisable personal data – are automatically detected and replaced with anonymous placeholders before any data is sent to the AI provider. This detection operates on all unstructured text inputs.

### 1.4.2 Layer 2 – Firm Name Exclusion

The firm's identity is stripped from all data before AI processing. The AI never sees the firm name. After the AI has generated its output, the firm name and relevant contextual details are re-inserted on SettlDay's own EU servers during post-processing. This means the AI model processes a pseudonymised payload with no knowledge of which firm it is analysing.

### 1.4.3 Layer 3 – Document Summarisation

When users upload documents (e.g., operational procedures, settlement reports), the full document content is never forwarded to the AI provider. Instead, documents are processed on SettlDay's EU servers:

1. Text is extracted from the uploaded document

2. The content is reduced to short, structured excerpts containing only operationally relevant information
3. PII within the excerpts is replaced with placeholders (Layer 1)
4. Only these sanitised excerpts are included in the AI prompt
5. Original document files are deleted after extraction

---

## 1.5 AI Provider Architecture

Property	Detail
<b>AI provider</b>	AWS Bedrock (managed AI inference service)
<b>EU entity</b>	Amazon Web Services EMEA SARL, Luxembourg
<b>Data residency</b>	EU regions only – no data leaves the EU
<b>Inference logging</b>	Disabled – inputs and outputs are not logged or stored by AWS after inference
<b>Data retention</b>	Zero – no customer data is retained by the AI provider
<b>Training use</b>	Never – AWS Bedrock contractual terms prohibit use of customer data for model training or improvement
<b>Anthropic's role</b>	None – Anthropic does not receive, process, see, or store any customer data. AI models run on AWS-managed infrastructure. The processor relationship is with AWS, not Anthropic. Anthropic is not a sub-processor.
<b>DPA</b>	AWS GDPR Data Processing Addendum with EU Standard Contractual Clauses in effect via AWS Customer Agreement

**Why Bedrock and not a direct AI provider API?** SettlDay uses AWS Bedrock specifically because it provides EU data residency, disabled inference logging, and a formal GDPR Data Processing Addendum. These guarantees are not available on the AI model provider's direct API. This is a deliberate architectural choice to meet the data residency and zero-retention requirements expected by regulated financial services clients.

---

## 1.6 MiFID II Considerations

All AI outputs are explicitly marked as **advisory – not investment advice**. SettlDay reports provide operational readiness assessments for T+1 settlement transition, not investment recommendations, trading strategies, or portfolio-specific advice.

Disclaimers on every report explicitly state that the content does not constitute investment advice and should not be relied upon as the sole basis for compliance decisions.

---

## 1.7 Record-Keeping and Auditability

For every AI generation event, SettlDay logs:

- **Model version** used for generation
- **Token usage** (input and output)
- **Cost** per generation stage
- **Duration** of processing
- **Timestamps** with full audit trail

These records are retained in SettlDay's EU database for auditability and regulatory compliance purposes.

---

## 1.8 Local Embeddings

The research assistant's document search capability uses **local inference** for embedding generation. Document embeddings are computed entirely on SettlDay's own EU infrastructure. No external API calls are made for embedding generation – no document content leaves SettlDay's systems during the indexing process.

---

## 1.9 Data Sent to the AI Provider

Data Category	Sensitivity	Pre-processing Applied
Firm type, firm size, region	Low (quasi-identifiers)	None required
Assessment category scores	Low (non-PII)	None required
Individual question answers	Low (operational maturity ratings)	None required
User questions (free text)	Low-Medium (could contain names or PII)	PII detection + firm name exclusion
Document excerpts	Low-Medium (summarised content)	Summarised on SettlDay servers; PII scrubbed

---

**No directly identifying personal data (names, emails, addresses) is sent to the AI provider. The firm name is never transmitted.**

---

## 1.10 Art. 22 – Automated Decision-Making

Criterion	Assessment
Automated decisions with legal or significant effect	<b>No</b> – all AI outputs are advisory reports for human review
Right to human intervention	Not applicable – no automated decisions are made
Profiling	<b>No</b> – assessment scoring is static and rule-based, not AI-driven

---

*For questions about AI governance, data minimisation, or the AI processing pipeline, contact [privacy@settlDay.com](mailto:privacy@settlDay.com).*

## 2. Data Flow Architecture

---

<b>Document ID</b>	TC-ARCH-001
<b>Classification</b>	Public – Trust Center

---

### 2.1 Quick Facts

Question	Answer
Where is my data stored?	EU only
Is my data used to train AI models?	Never
Is my firm name sent to the AI?	Never
How long does the AI provider retain my data?	Zero – inference logging is disabled on AWS Bedrock
Which AI provider handles my data?	AWS Bedrock (EU regions) – not Anthropic directly
Can I delete my data?	Yes – self-service deletion, hard delete within 30 days

---

### 2.2 Data Flow Steps

#### 2.2.1 Step 1 – Your Assessment Data

Assessment responses, firm profile data, and uploaded documents are transmitted over TLS and stored encrypted (AES-256) on EU infrastructure.

##### What is stored at this stage:

- Firm type, size, and jurisdictions of operation
  - Assessment question responses and maturity scores
  - Operational KPIs (fail rates, STP rates) if provided
  - Free-text comments (if entered)
  - Uploaded document content (temporary – see Step 2)
-

## 2.2.2 Step 2 – Data Minimisation and Pseudonymisation

Before any data is sent to an AI system, Settlday applies the following transformations entirely on its own EU infrastructure:

Transformation	Detail
<b>Firm name removed</b>	Your firm's name is stripped from all data before AI processing and never transmitted to the AI provider
<b>Documents summarised</b>	Uploaded documents are reduced to short structured excerpts on Settlday servers; original files are deleted after extraction
<b>PII replaced</b>	Names, email addresses, and other personal identifiers in free-text fields and document excerpts are replaced with anonymous placeholders
<b>Minimum data forwarded</b>	Only the data fields required to generate the report are included – no identifiers, no account metadata

The result is a pseudonymised, firm-name-free payload containing only what is necessary for the AI to generate a relevant report.

## 2.2.3 Step 3 – AI Processing

The pseudonymised payload is sent to **AWS Bedrock**, Amazon's managed AI inference service, operating exclusively in EU regions.

Property	Detail
<b>Infrastructure owner</b>	Amazon Web Services EMEA SARL (EU entity, Luxembourg)
<b>Data residency</b>	EU regions only – no data leaves the EU
<b>Inference logging</b>	Disabled – inputs and outputs are not logged or stored by AWS
<b>Training use</b>	Never – AWS Bedrock contractual terms prohibit use of customer data for model training or improvement
<b>Anthropic's role</b>	None – Anthropic does not receive, process, see, or store any customer data when models run on Bedrock. The processor relationship is with AWS, not Anthropic

**Why Bedrock?** Settlday uses AWS Bedrock specifically because it provides EU data residency, disabled inference logging, and a formal GDPR Data Processing Addendum – guarantees that are not available on the AI model provider's direct API. This is a deliberate architectural choice to meet EU data residency and zero-retention requirements.

## 2.2.4 Step 4 – Your Report

The raw AI output is returned to SettlDay's EU servers. Post-processing on SettlDay infrastructure:

- Re-inserts your firm name and relevant contextual details
- Validates and sanitises the AI output
- Structures the report into sections (PDF)
- Stores the final report encrypted on SettlDay's EU infrastructure

Your report is only ever assembled in complete form on SettlDay's own servers. The AI never sees your firm name or personal details – they are added only after the AI has finished processing.

## 2.3 Compliance Framework

Obligation	Status	Detail
GDPR Art. 30 – ROPA	Maintained	Full records of processing activities documented
Data Processing Agreement	Available	DPA with EU Standard Contractual Clauses available for enterprise customers
DPIA	Completed	Data Protection Impact Assessment completed and signed
Self-service data deletion	Implemented	Account deletion with grace period followed by hard delete of all personal data
Data portability	Implemented	Full JSON export of all personal data available
EU AI Act – transparency	Compliant	AI-assisted report generation disclosed to users; no high-risk AI system classification
Breach notification	72-hour commitment	Supervisory authority notified within 72 hours of discovery (GDPR Art. 33)
Data subject rights	Implemented	Art. 15-22 procedures documented and self-service where possible

## 2.4 Data Flow Summary

Stage	What happens	Where	Data exposed to AI?
1. Collection	Assessment responses and documents submitted via TLS	SettlDay EU	No
		SettlDay EU	No

Stage	What happens	Where	Data exposed to AI?
<b>2. Pre-processing</b>	Firm name removed, PII replaced with placeholders, documents summarised		
<b>3. AI inference</b>	Pseudonymised payload sent to AWS Bedrock	AWS Bedrock EU	Only anonymised data – no PII, no firm name
<b>4. Post-processing</b>	Firm name re-inserted, report validated and rendered as PDF	SettlDay EU	No
<b>5. Delivery</b>	Encrypted report stored and delivered to user	SettlDay EU	No

## 2.5 EU Data Residency

All data remains within the EU at every stage of processing:

Data Category	Location	Leaves EU?
Firm name and identity	SettlDay EU servers only	No – never transmitted to any third party
User personal data (name, email, credentials)	SettlDay EU infrastructure	No
Raw assessment responses	SettlDay EU database (encrypted)	No
Uploaded documents	SettlDay EU servers (deleted after summarisation)	No
Generated reports	SettlDay EU database (encrypted)	No
Pseudonymised AI payload (no PII, no firm name)	SettlDay EU servers and AWS Bedrock EU	No – stays within EU AWS regions
Payment data	Stripe (US, EU SCCs in effect)	Card data handled by Stripe.js – never touches SettlDay servers

Anthropic has zero data access. Claude runs on AWS Bedrock managed infrastructure; the processor relationship is with AWS.

The full Sub-Processor Register is included in the Vendor Due Diligence Pack, available at the [SettlDay Trust Center](#). For questions about data flows, contact [privacy@settlday.com](mailto:privacy@settlday.com).