SPONTAINE

Privacy, Governance & Security Whitepaper

SPONTAINE

spontaine.com
Intuon Analytics Pvt Ltd

01

Governance By Architecture

Enterprise software platforms that process sensitive business data carry an implicit obligation.
The architecture itself must enforce trust, not merely promise it.
At Spontaine, more than merely an aspiration, it is the founding principle behind every layer of the platform.

The Challenge
Organisations adopting AI-powered analytics face a new and expanded governance surface. Legacy controls like access lists, export policies, manual anonymisation were not designed for systems where an AI model constructs queries, synthesises answers, and creates output in natural language.

Spontaine addresses this by treating governance as a first-class architectural concern.
The platform is built on four layers :
Data Ingestion, Semantic Intelligence, Generative Interface, and Governed Operations.
Privacy, auditability, and access control are enforced at each transition. The result is a system where compliance is structural, not procedural.

This document sets out how Spontaine operationalises that commitment through the AI pipeline design, the deployment model, the contractual stack, and the flexible data residency architecture , giving procurement teams, DPOs, and CXOs a complete and accurate picture.

02

How Your Data Moves Through Spontaine

The Enterprise Data Platform

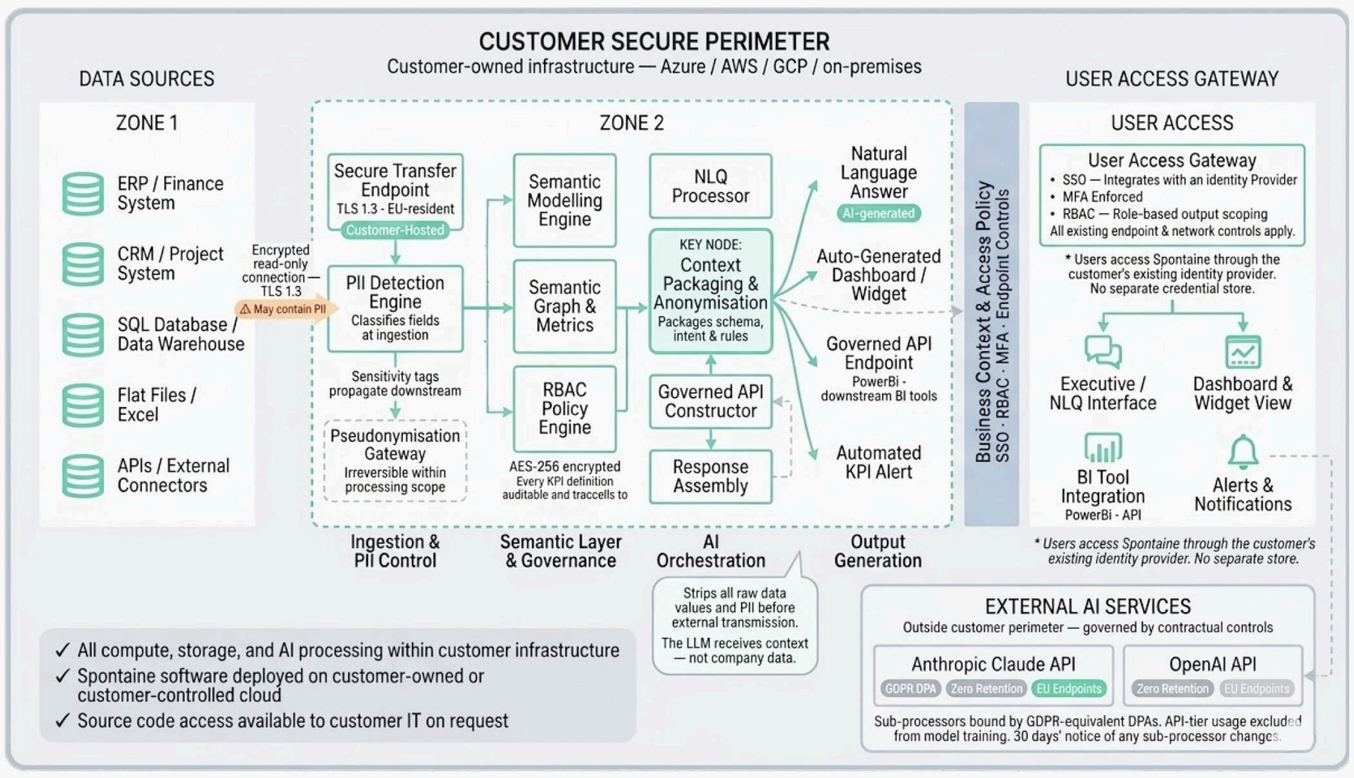

Figure 1: Spontaine's four-layer production architecture. All compute, storage, and AI processing occur within the EU Data Residency Zone. The LLM receives only sanitised context - never raw client data.

The AI Pipeline: Models Never See Client Data Directly

Spontaine's AI capabilities are delivered through a native orchestration pipeline that only communicates over governed, context-aware APIs.
The sequence is critical from a privacy standpoint:

- *User submits a natural-language query*
- *The internal workflow layer enriches and sanitises the request*
- *A PII anonymisation layer strips or pseudonymises any sensitive data - this is a technical control enforced in code, not a configuration option*
- *The LangGraph orchestrator resolves intent and selects the appropriate internal tool*
- *Spontaine APIs fetch the relevant data from its internal semantic engine*
- *Context, structure, and intent are sent to the model, which constructs a governed API request - on full guardrails.*
- *The governed API executes against the semantic layer; the result is returned to the user*

Control	PoC Environment	Production Environment
Encryption in transit	TLS 1.2+	TLS 1.2+
Encryption at rest	AES-256	AES-256
Access control	RBAC - granular, designated users	Standard enterprise security controls.
Audit logs	Immutable;	Immutable;
PII handling	Anonymised before every LLM call	Anonymised before every LLM call
Penetration testing (VAPT)	Not standard for PoC scope	Ongoing;
Data residency	EU-only (Spontaine-managed)	Customer-controlled infrastructure
MFA	Not standard for PoC scope	Enforced; integrated with customer IdP

03

Governance

In all Spontaine engagements, Spontaine acts as Data Processor and the customer acts as Data Controller. Anthropic, OpenAI, Azure, AWS, and DigitalOcean serve as sub-processors. A standard GDPR Article 28 Data Processing Agreement is available at spontaine.com/dpa and forms part of every commercial engagement.

Roles & GDPR

- Article 28 DPA in place for all EU client engagements
- Article 30 Records of Processing Activities (RoPA) maintained by Privacy Officer
- Article 32 Technical and organisational measures documented and implemented
- Article 33 Breach notification to supervisory authority within 72 hours
- Contractual obligation: 24-hour breach notification to the customer — stricter than the GDPR minimum
- Data subject rights procedures in place: access, erasure, portability, rectification
- Data Protection Impact Assessment (DPIA) available upon request
- Sub-processor register maintained; 30 days' written notice of changes

Data Minimisation in the AI Pipeline

Spontaine's approach to data minimisation goes beyond restricting what data is ingested. The AI pipeline enforces minimisation at the point of use: only the metadata, schema context, and anonymised structural signals required to construct a query are transmitted to the LLM. Raw field values , and especially any PII fields identified at ingestion, are never included in LLM calls.
This operationalises Article 5(1)(c) of the GDPR at an architectural level.

Explainability and Reasoning Chains

Spontaine AI outputs include transparent reasoning chains. The system discloses data sources consulted, the logic applied, and any limitations in the underlying data.
The screenshot below illustrates a live Spontaine interaction in which the AI assistant surfaces its reasoning process alongside its answer.

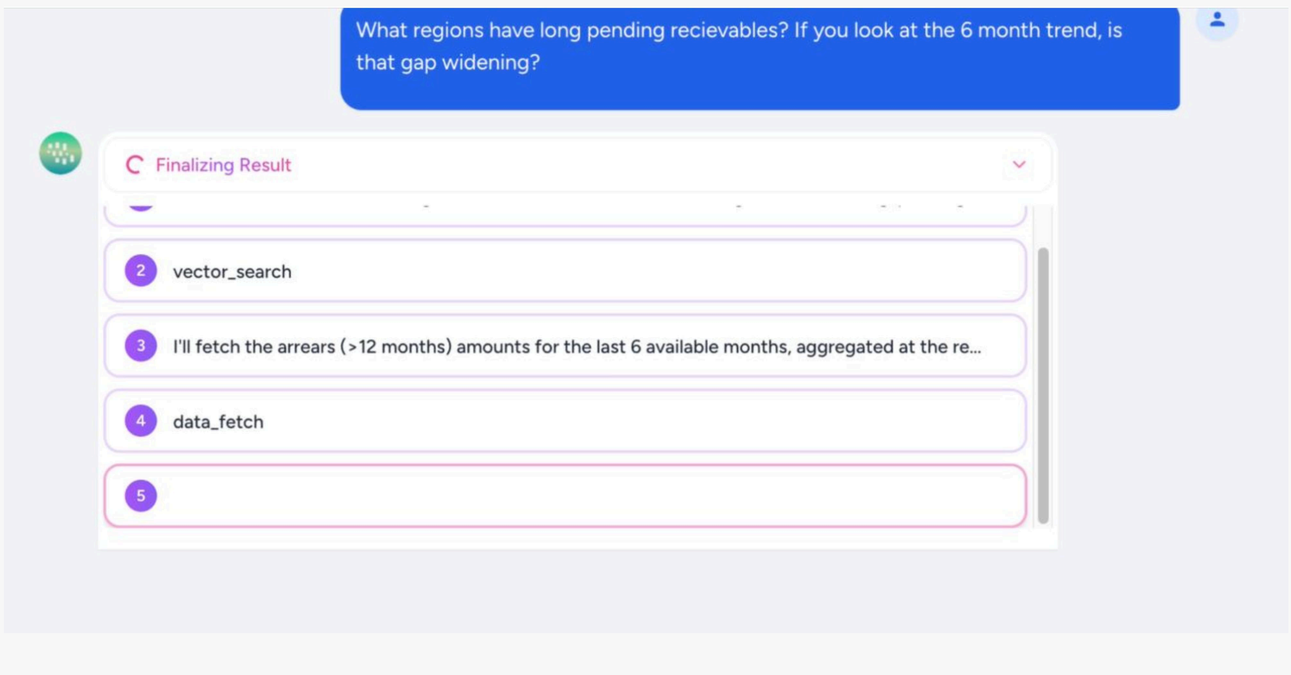

Figure 2: Spontaine's AI assistant discloses its reasoning chain and data sources alongside every generated answer. This transparency is non-suppressible and supports both internal audit requirements and the EU AI Act's transparency obligations.

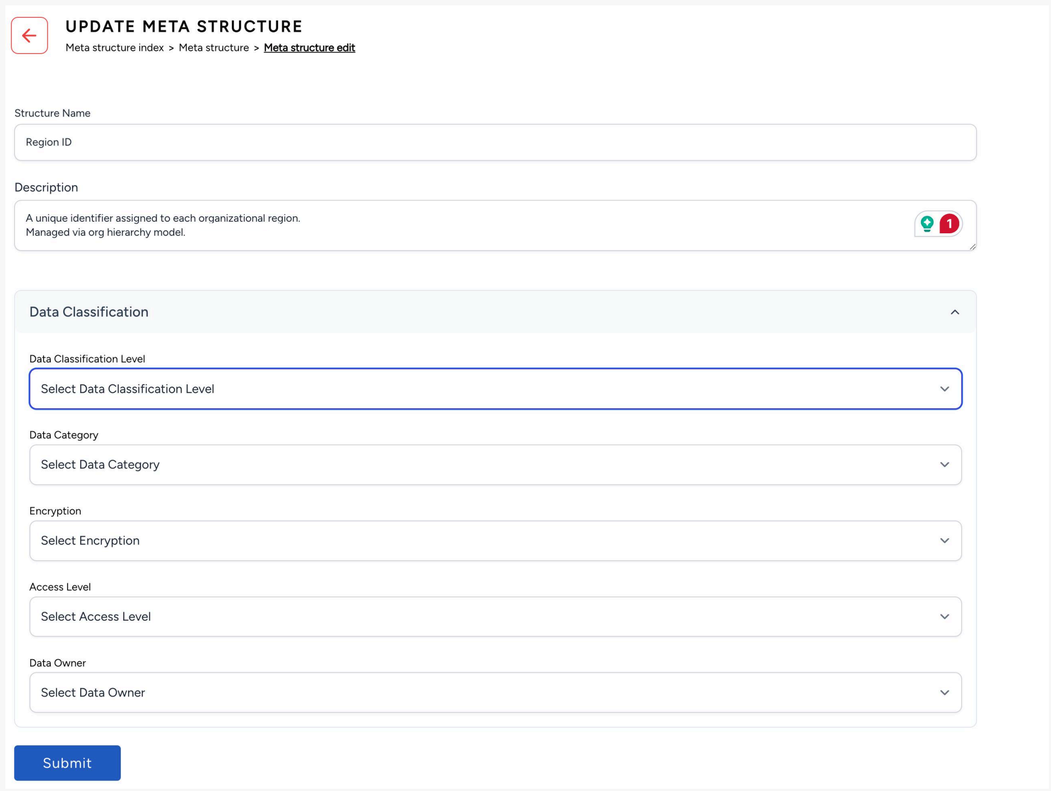

Figure 3: The PII anonymisation layer is subject to automated testing within the CI/CD pipeline and is reviewed at every production release. Any failure is classified as a critical security incident requiring immediate escalation to the Information Security Officer.

04

Compliance Posture

Data Governance

Spontaine maintains a formal sub-processor register. All sub-processors are bound by GDPR-equivalent contractual obligations. Customers receive 30 days' written notice of any changes.

Sub-Processor	Processing Purpose	Location	Data Retention
Microsoft Azure	EU PoC & production hosting	EU (NL / IE)	Per customer config
Amazon AWS	EU PoC hosting	EU (DE / IE)	Per customer config
DigitalOcean	EU PoC & dev hosting	EU (NL)	Per customer config
Gemini API	AI inference — context only, no raw data	EU endpoints	Zero (API tier)
OpenAI API	AI inference — context only, no raw data	EU endpoints	Zero (API tier)
GitHub	Source control — no client data	Global (code only)	N/A

Spontaine's governance architecture is not a compliance layer added to a data product - it is the mechanism by which the product delivers trustworthy intelligence.
The PII anonymisation pipeline, the Semantic Layer's auditable lineage, the client-hosted production deployment model, and the contractual EU data residency commitments are all expressions of the same principle: governed intelligence means the platform enforces trust at every layer.

Framework / Certification	Status	Owner / Reference
ISO 27001:2022	Certified - continuous monitoring via Enterprise compliance platform	COO/ISO
GDPR (EU) 2016/679	Compliant - DPA, RoPA, breach procedures in place	spontaine.com/dpa
EU AI Act 2024/1689	Limited Risk - transparency obligations met	AI AUP: ISMS-AI-001
SCCs (2021/914)	In place for India - EU access pathway	Included in client DPAs
Azure / AWS / DigitalOcean	Certified: ISO 27001, SOC 2, PCI DSS	Infrastructure partners

SPONTAINE

spontaine.com

For enterprise procurement teams and DPOs, the relevant documentation is available immediately:

- *Standard DPA: https://spontaine.com/dpa*
- *DPIA available on request*
- *Compliance monitoring report (point-in-time) available on request*
- *Client references available on request*

Intuon Analytics Private Limited *2026*

Contributors

Swaraj Krishnan

CTO

(System Architecture)

At Spontaine, Swaraj leads the evolution of the core architecture, from the proprietary AI Semantic Layer to the future-proof Infinity Ingestion Platform.

swaraj@intuonfx.com

Radhika Ravindran

COO

(Governance, Delivery)

At Spontaine, Radhika owns the entire customer journey from onboarding to operational excellence, guaranteeing a world-class experience that drives adoption and revenue expansion.

radhika@intuonfx.com

For security review requests, sub-processor queries, or to discuss the EU data residency architecture in detail, contact: radhika@intuonfx.com or desk@spontaine.com.