A P R I T Y

Aprity — AI & LLM Data Handling

AI & Large Language Model Data Handling

PUBLIC

Version 1.0 | March 2026 | aprity.ai

1. Executive Summary

Aprity uses Large Language Models (LLMs) to **explain and document** Salesforce metadata. This document answers the most common questions about how your data interacts with AI providers.

Key guarantees:

- Only **metadata** is sent to AI (code structure, field names, automation logic) — never business data
- Every AI interaction is **logged and auditable** for 30 days
- You can **bring your own LLM** (BYO-LLM) for full control over data residency
- AI is used **only for explanation** — all business logic is deterministic

2. What Is Sent to AI Providers?

2.1 Data Included in AI Prompts

Data Type	Examples	Why?
Apex code	Class definitions, trigger logic, test methods	To explain what the code does
Flow definitions	Flow steps, decision criteria, actions	To document automation
Object schema	Field names, data types, relationships	To describe the data model
Validation rules	Rule formulas, error messages	To document business rules
Permission config	Permission set names, field-level security	To describe security model
Dependency graphs	Which components call which	To explain execution flow

2.2 Data Never Sent to AI

Data Type	Guarantee
Business records	No Account, Contact, or Opportunity data — ever
User data	No names, emails, phone numbers, or PII
Passwords & tokens	No credentials of any kind
File contents	No attachments, documents, or images
Session data	No Salesforce session IDs or access tokens

> ***Technical guarantee**: Aprity uses the Salesforce **Metadata API**, which cannot access record-level data. It is structurally impossible for business data to reach an AI provider.*

3. How AI Is Used

3.1 Deterministic First, AI Second

Aprity follows a strict principle: **if a computation can be done deterministically, AI is not used**.

Task	Method	AI Involved?
Dependency resolution	Graph traversal algorithm	No
DML classification	Deterministic BFS on execution graph	No
Security rule analysis	Permission matrix computation	No
Code explanation	LLM analysis	Yes
Business rule documentation	LLM synthesis	Yes
Process flow description	LLM narration	Yes

3.2 AI Outputs

AI produces:

- **Functional documentation**: Business-friendly explanations of what code does
- **Technical documentation**: Developer-oriented analysis of code structure
- **Business rule summaries**: Plain-language descriptions of validation and automation logic
- **Impact analysis**: Cross-component dependency explanations

AI does **not** make decisions about:

- Scan orchestration
- Data security or access control
- Metadata classification
- Dependency resolution

4. AI Provider Options

4.1 Supported Providers

Provider	Auth Method	Data Location	Default?
Anthropic Claude	API Key	US	Yes (configurable)
Azure OpenAI	Azure Managed Identity	Your Azure region	Optional
OpenAI Direct	API Key	US	Optional
Xai (Grok)	API Key	US	Optional
Custom / Self-hosted	API Key or Bearer	Your infrastructure	Optional

4.2 BYO-LLM (Bring Your Own LLM)

Enterprise customers can configure their own AI endpoint:

- **Azure OpenAI in your subscription**: Data stays in your chosen Azure region (e.g., France Central, West Europe)

- **Self-hosted models**: Full control — data never leaves your infrastructure
- **Any OpenAI-compatible endpoint**: Custom API with your own terms

Benefits of BYO-LLM:

- Data residency in your chosen jurisdiction
- Your own data processing agreement with the AI provider
- No third-party data exposure
- Full control over model selection and versioning

4.3 Provider Switching

Aprity supports **zero-downtime provider switching**:

- Configure multiple endpoints with priority ordering
- Automatic failover if primary provider is unavailable
- Per-tenant override: different tenants can use different providers

5. AI Data Residency

5.1 Default Configuration

Component	Location
Aprity platform	Azure France Central (EU)
AI prompts in transit	Encrypted via TLS 1.2+
AI processing (default)	Depends on provider (see table above)
AI interaction logs	Azure France Central (EU)

5.2 EU-Only Configuration

To ensure all data stays within the EU:

1. Deploy **Azure OpenAI** in a European Azure region
2. Configure it as the primary endpoint in Aprity
3. Result: metadata never leaves the EU, including AI processing

6. AI Provider Data Policies

6.1 No Training on Your Data

Aprity uses AI providers via their **API-only** agreements:

Provider	Training on API data?	Data retention?
Anthropic Claude (API)	No	Zero retention (API terms)
Azure OpenAI	No	No (with opt-out)
OpenAI (API)	No (with opt-out)	Zero retention (API terms)

> *These are standard API terms from the respective providers. Aprity does not use consumer-grade AI interfaces.*

6.2 No Data Sharing Between Tenants

- Each AI request is **independent** — no shared context between tenants
- No fine-tuning or model adaptation using customer data
- No aggregation of data across tenants

7. Audit Trail

7.1 What Is Logged

Every AI interaction is logged with:

Field	Description
Tenant ID	Which customer initiated the analysis
Scan ID	Which scan the analysis belongs to
Provider	Which AI provider was used
Model	Specific model version (e.g., Claude Sonnet 4)
Token count	Input and output token usage
Duration	Processing time
Status	Success, parse error, or failure
Full prompt	System prompt + user message (stored in Blob Storage)
Full response	AI provider response (stored in Blob Storage)

7.2 Retention

- **Metadata** (provider, tokens, duration): 30 days (automatic deletion)
- **Content** (prompts, responses): 30 days in Blob Storage (automatic deletion)
- After 30 days, all AI interaction data is **permanently deleted**

7.3 Access

- AI logs are **tenant-scoped** — only accessible by the tenant that generated them
- Used for: quality assurance, cost attribution, debugging, and audit compliance
- Available to the customer upon request

8. Prompt Security

8.1 Prompt Construction

- System prompts are **template-based** — no customer data in system prompts
- User prompts contain only metadata (code, schema, config names)
- Prompts are constructed server-side — no user input injection possible

8.2 Response Handling

- AI responses are parsed as **structured JSON** — not executed as code
- Invalid or unparseable responses are logged and discarded
- AI output is used for documentation only — never for access control or business decisions

9. Frequently Asked Questions

Q: Can Aprity read my Salesforce records (Accounts, Contacts, etc.)?

A: No. Aprity uses the Metadata API, which provides access to code and configuration only. Record-level data is inaccessible.

Q: Is my data used to train AI models?

A: No. All AI providers are used via API agreements that explicitly exclude training on customer data.

Q: Can I use my own AI provider?

A: Yes. Aprity supports BYO-LLM — configure your own Azure OpenAI, self-hosted model, or any compatible endpoint.

Q: How long is my data retained by the AI provider?

A: Zero retention under standard API terms. Aprity's own logs are retained for 30 days, then automatically deleted.

Q: Can another customer see my AI analysis?

A: No. All AI interactions are strictly tenant-scoped with partition-key isolation at the database level.

Q: What happens if the AI provider is down?

A: Aprity supports automatic failover to a secondary provider. If no provider is available, the scan pauses and retries automatically.

Contact

For questions about AI data handling:

Email: security@aprity.ai

Website: https://aprity.ai

This document describes Aprity's AI data handling practices as of March 2026. For the most current information, please contact our security team.