DATA PROTECTION REPORT

AI Data Privacy and Security

Prepared for: **Example Ltd (sample data)**

Audit Reference: **JNVOY-4A8D80FD66DF**

Generated: **26 October 2026 at 15:27 UTC**

Period: **2026-09-01 to 2026-10-01**

Prepared by Jnvoy

jnvoy.com | github.com/Jeevan2191999/jnvoy

1. Statement

No PII detected by Jnvoy was transmitted to any LLM API during the audit period.

All requests routed through Jnvoy during the audit period were inspected for personally identifiable information (PII) before being forwarded to any large language model. Sensitive data identified in any request was redacted and replaced with reversible tokens prior to transmission. The original values were restored in the response before delivery to the end user. Detection is automated and is not guaranteed to be exhaustive. Data that does not match a configured detector, such as unusual formats or free-text identifiers, may not be redacted. This report documents the controls applied and their results. It is not a guarantee that no personal data was ever transmitted.

Statement	Value
Audit Reference	JNVOY-4A8D80FD66DF
Company	Example Ltd (sample data)
Audit Period Start	2026-09-01
Audit Period End	2026-10-01
Report Generated	2026-10-26 15:27 UTC
Detected values sent unredacted	NONE

2. Executive Summary

During the audit period, Jnvoy processed **1,250** API requests on behalf of Example Ltd (sample data). Of these requests, **847** instances of personally identifiable information were detected and redacted before any data was forwarded to an external LLM provider. The semantic cache resolved **34.2%** of requests without contacting any external provider, reducing both cost and data exposure.

Metric	Value	Notes
Total API calls processed	1,250	All requests inspected for PII
PII instances detected and redacted	847	Replaced with tokens before the text was sent
Cache hit rate	34.2%	Requests resolved without external call
Average request latency	1824ms	Including PII detection overhead
Detected values sent unredacted	None	Values it misses are not counted

3. PII Detection Breakdown

The following categories of personally identifiable information were detected and redacted during the audit period. All instances were replaced with reversible tokens before the text was sent and restored in the response. This report counts only what Jnvoy detected. It cannot show that no other personal data was present.

PII Type	Instances Detected
Email Address	312
Person	289
Phone Number	156
Uk Postcode	54
Ni Number	36
TOTAL	847

4. Technical Methodology

PII Detection Engine

Jnvoy uses Microsoft Presidio, an enterprise-grade NLP-based PII detection engine, combined with pattern-based recognisers for UK-specific data types. Detection runs locally on Jnvoy infrastructure. No data is sent to any external service for the purpose of PII detection. The following entity types are supported: person names, email addresses, phone numbers, credit card numbers, IBAN codes, IP addresses, UK NHS numbers, National Insurance numbers, UK postcodes, sort codes, account numbers, passport numbers, API keys, and secret tokens.

Tokenisation and Restoration

When PII is detected, each instance is replaced with a unique reversible token before the request is forwarded to the LLM. The token map is held in memory for the duration of the request and is used to restore original values in the LLM response before delivery to the end user. The token map is never persisted. Only PII type counts and timestamps are written to the audit log.

Audit Logging

Every request is recorded in a PostgreSQL audit database with a unique audit ID, timestamp, PII type counts, provider used, model used, latency, and cache status. Raw PII values are never written to the audit log. Only type counts are stored.