eFraud Investigator

AI Data Privacy & Confidentiality

How Google Vertex AI Handles Your Clients' Financial Documents • A Technical Reference for Forensic Accountants & Investigators

Platform: Google Vertex AI (Google Cloud enterprise AI)
Governing terms: Google Cloud Platform Terms of Service, Cloud Data Processing Addendum, and Service Specific Terms
Version: 2.0 • June 2026

> **ℹ Executive summary for forensic accountants**
>
> When you upload a bank statement, brokerage record, or financial PDF to eFraud Investigator, it is processed by Google Vertex AI — Google Cloud's enterprise AI platform, not the consumer Gemini app. Under Google Cloud's contractual terms, your clients' data is not used to train AI models and is encrypted in transit and at rest. As eFraud's Google Cloud project is configured — invoiced billing, data caching disabled, no search grounding — Google retains no copy of your clients' document content after processing completes. This document explains how, and points to Google's published terms.

1. The Platform: Vertex AI, Not Consumer Gemini

Forensic accountants often hear "Gemini" and picture the consumer chatbot at gemini.google.com. eFraud Investigator does not use that product. It uses Google Vertex AI, the enterprise AI platform on Google Cloud, governed by the Google Cloud Platform Terms of Service, the Cloud Data Processing Addendum (CDPA), and the Service Specific Terms. Understanding the difference is essential.

Consumer Gemini app (gemini.google.com)	Google Vertex AI (what eFraud uses)
General-public product	Enterprise AI platform on Google Cloud
Subject to consumer terms of service	Subject to GCP Terms, the CDPA, and the Service Specific Terms
Inputs may be reviewed to improve products	Training restriction is contractual (Section 17, Service Specific Terms)
No enterprise compliance posture	Runs on Google Cloud's independently audited infrastructure (§6)
Shared consumer environment	Project-level isolation; data not shared across customers

2. What Happens to Your Client's Document

The following describes what occurs from the moment you initiate processing to the moment you receive the extracted data.

Step	What happens	Privacy safeguard
1. Upload	You upload the PDF within eFraud Investigator's secure interface.	Encrypted in transit; stored in your eFraud account only.
2. Pre-scan	The extraction engine sends page images to Vertex AI to identify statement type and structure.	Transmitted over encrypted HTTPS; processed in Google's enterprise environment.
3. Extraction	Vertex AI reads each page and returns structured transaction data as JSON.	Caching disabled and invoiced billing in effect — no copy retained by Google; data not used for training.
4. Return	The JSON is received by eFraud's servers and written to your database.	Stored in eFraud's isolated, encrypted database, not Google's systems.
5. Completion	Processing ends; no document content remains on Google's side.	Zero data retention — no copy persists with Google.

3. Zero Data Retention: How It Is Achieved

Vertex AI retains customer data only in specific, limited scenarios — each of which can be controlled. eFraud's project is configured so that none of them retains your clients' document content. The table below maps each retention vector to eFraud's configuration.

Retention vector	Default Vertex AI behavior	eFraud's configuration
Data caching	Inputs and outputs cached up to 24 hours at the project level.	Disabled at the project level (disableCache = true) — no caching retention.
Abuse-monitoring prompt logging	Prompts may be logged up to 90 days — but only for accounts without invoiced billing.	Not applicable: eFraud uses an Invoiced Cloud Billing account, which is out of scope for abuse-monitoring logging.
Grounding with Google Search	Stores prompts and output for 30 days; cannot be disabled.	Not used — eFraud's pipeline makes no search-grounding calls.
Live API session resumption	Off by default.	Not used — not enabled in eFraud's integration.
Model training	Not used without permission (Section 17).	Not used — protected by Section 17.

Retention vector	Default Vertex AI behavior	eFraud's configuration
Extracted JSON output	Returned to caller, not stored by Google.	Received by eFraud, stored in eFraud's database only.

What this means

As configured — invoiced billing, caching disabled, no search grounding, no Live API — Google retains no copy of your client's document content after the API response is returned. A subpoena served on Google for that content would have nothing responsive to produce. eFraud's database is the system of record.

Configuration note

These statements describe eFraud's current Vertex AI configuration, verified June 2026. The configuration is under eFraud's control and can be re-verified on request. eFraud will update this document if its configuration changes.

4. Google's Contractual Commitments

The most important commitment for evidentiary and professional-responsibility purposes is the training restriction. Under Section 17 ("Training Restriction") of Google's Service Specific Terms, Google commits by contract not to use customer data to train or fine-tune any AI/ML models without the customer's prior permission or instruction. This applies to all managed models on Vertex AI, including general-availability and pre-GA models.

This is not a policy statement — it is a binding contractual obligation under the Google Cloud Platform Terms. Google's data handling is further governed by the Cloud Data Processing Addendum (CDPA), under which customers control their data, control where and how it is processed, and have it processed only in accordance with the customer's instructions.

5. Encryption — In Transit and At Rest

All financial data is protected by industry-standard encryption at every stage of the pipeline.

Encryption layer	Standard applied
Data in transit (page images sent to Vertex AI)	TLS 1.2+ (transport-layer security)
Data at rest (if temporarily held in Google's infrastructure)	AES-256, Google-managed keys
Customer-Managed Encryption Keys (CMEK)	Available for additional control — keys held by eFraud, not Google
Data in transit (eFraud to your browser)	TLS 1.2+, enforced at the CloudFront layer

Encryption layer	Standard applied
Data at rest (extracted records in eFraud's database)	AWS RDS encryption, AES-256

No unencrypted financial data traverses the network at any point in the eFraud Investigator pipeline.

6. Google Cloud's Infrastructure Certifications

The Google Cloud infrastructure underpinning Vertex AI — the platform that processes your documents — maintains the following compliance certifications, verified by independent third-party auditors.

Certification / standard	Relevance to financial investigations
SOC 2 Type II	Validates ongoing controls for security, availability, and confidentiality — the standard most law and CPA firms require of cloud vendors.
ISO/IEC 27001	International information-security management standard.
ISO/IEC 27017 / 27018	Cloud-specific security controls and protection of personally identifiable information.
HIPAA (via BAA)	Infrastructure-level controls sufficient for highly regulated sensitive data — analogous to PII in financial records.
FedRAMP (select services)	US federal government authorization — a rigorous data-handling standard.

> ⚠ **An important distinction**
>
> These certifications are held by Google for the Vertex AI infrastructure. They describe the platform eFraud Investigator builds on; they are not certifications held by eFraud Services Inc. eFraud does not represent that it holds SOC 2 or any other certification in its own name.

7. Data Isolation — No Cross-Customer Exposure

Google Cloud enforces strict multi-tenancy isolation, so other Vertex AI users cannot access your data through shared infrastructure.

- Each eFraud API call executes within eFraud's dedicated Google Cloud project context.
- Project-level privacy is enforced; no data crosses project boundaries.
- VPC Service Controls can be configured to prevent data exfiltration across organizational boundaries.
- Google does not aggregate or combine customer data across accounts.

Your client's statement cannot be seen by another eFraud customer, and it cannot be accessed by other organizations using Vertex AI for unrelated purposes.

8. Where Your Data Lives After Extraction

Once Vertex AI returns the extracted transaction data, Google has no further involvement. The extracted data resides in:

- Database — eFraud's AWS RDS (us-east-1), encrypted at rest.
- Case isolation — separated per case, so each investigation is logically distinct.
- Access control — governed by your eFraud account credentials and role permissions.
- Tenancy — accessible only to your authorized users, never shared with other eFraud customers.

The original PDF is not retained by eFraud's AI extraction service after processing completes; eFraud stores the structured extraction output.

9. Data Residency

eFraud's pipeline currently uses Vertex AI's "global" endpoint, which routes requests dynamically across Google's regions for availability and access to current models. Because of this, eFraud does not represent that processing occurs in any specific country or region. For engagements that require regional data residency, Vertex AI supports pinning requests to a specific region, which eFraud can configure on request.

10. Professional Responsibility Considerations

Forensic accountants and fraud investigators working under professional standards (AICPA, ACFE, and — for attorneys — bar rules) have obligations to maintain client confidentiality when using third-party tools. The table below addresses those obligations; it describes how the platform supports them and is not legal advice.

Professional concern	How eFraud Investigator supports it
Client-data confidentiality	Processing on an enterprise platform under contract, with a no-training restriction and no vendor retention, supports the reasonable-measures standard for third-party tools.
Privilege (white-collar defense)	Data is processed by an enterprise platform under contract and is not retained by the vendor. This supports maintaining confidentiality; counsel should assess privilege implications for a given matter.
Chain of custody / evidence integrity	Vertex AI extracts data; it does not modify the source. The original PDF remains the evidentiary artifact. See eFraud's Methodology for Financial Data Extraction & Verification.

Professional concern	How eFraud Investigator supports it
Government investigations (subpoena)	Extracted data resides in eFraud's database under your account. With zero retention, Google holds no copy of the client's document content.
Engagement-letter disclosure	Consider disclosing that document extraction is performed by an AI tool operated on enterprise cloud infrastructure with contractual data protections.

11. Questions You Can Expect From Clients

Q: Does Google read my bank statements?
A: Vertex AI processes the page images to extract transaction data — analogous to OCR. With caching disabled and invoiced billing, Google retains no copy of the content, does not use it for model training, and has no ongoing access after the API call completes.

Q: Could Google employees see my client's statements?
A: Abuse-monitoring prompt logging is the only path by which Google would retain prompt content, and it does not apply to invoiced billing accounts such as eFraud's. Your client's documents are not used to train AI models.

Q: Is using an AI tool like this consistent with my professional obligations?
A: It is designed to be — provided the tool uses an enterprise, contractually bound platform (not a consumer app), encrypts data in transit and at rest, does not retain client data, and does not share data across customers. eFraud meets these conditions through Google Vertex AI. You should confirm your own obligations with reference to your governing standards.

Q: What if I receive a subpoena for my client's data?
A: The extracted records reside in your eFraud account database. With zero retention, Google holds no copy of the client's document content, so a subpoena served on Google would return nothing responsive. eFraud's database is the system of record.

12. Reference Sources

The following Google Cloud documentation and commitments underpin the statements in this document:

- Google Cloud Service Specific Terms, Section 17 "Training Restriction" — cloud.google.com/terms/service-terms
- Vertex AI Generative AI Data Governance — cloud.google.com/vertex-ai/generative-ai/docs/data-governance
- Vertex AI Abuse Monitoring — cloud.google.com/vertex-ai/generative-ai/docs/learn/abuse-monitoring
- Google Cloud Data Processing Addendum (CDPA) — cloud.google.com/terms/data-processing-addendum

- Google Cloud AI/ML Privacy Commitment — cloud.google.com/terms/aiml-privacy-commitment
- Google Cloud HIPAA Compliance — cloud.google.com/security/compliance/hipaa-compliance
- Google Cloud Compliance Certifications — cloud.google.com/security/compliance

ℹ About this document

This document describes eFraud Investigator's current data-handling configuration on Google Vertex AI and is provided for informational purposes. It is not legal advice. Professionals should review their own confidentiality and privilege obligations, and any required client disclosures, with reference to their governing standards and, where appropriate, counsel.

eFraud Services Inc. • Naples, FL • efraudservices.com • June 2026