Security overview

The five questions every security review asks us, answered before you ask. This document needs no NDA. The **SOC 2 Type II report**, penetration test summary, subprocessor list with locations, and Data Processing Addendum are available under NDA, and we supply the mutual NDA ourselves.

1 What the agents access, and least privilege

We read the accounts receivable data needed to work the invoices you place with us, and nothing else. Connections are read-only wherever the source system supports it. Where you enable write-back, it is scoped field by field with your team and stays off until you sign off on that field list.

Money never routes through AgentCollect. Settlement goes to your merchant account and your customers pay you directly, so no credential we hold has a payout destination we control. The worst case of a compromised read credential is data exposure, not an unauthorized transaction.

SOURCE	WHAT WE READ	POSTURE
ERP and accounting	Open invoices, customer master, payment status.	`Read only` by default
Bank feed	Transactions, to match payments to invoices.	`Read only` no payment initiation
Payment gateway	Payment status on placed invoices.	Scope depends on the vendor, disclosed per integration

Some vendors issue credentials with no scope picker, meaning the credential is broader than the single call we make. Where that is true for one of your systems, we tell you which vendor, what we actually call, and how to revoke us without involving us. We put that in writing during the build rather than leaving you to discover it.

2 Where your data lives, and for how long

Stored and processed in **AWS us-east-1** for US engagements. **AES-256** at rest with keys managed and rotated in AWS KMS, **TLS 1.3** in transit. Internal access is named-individual only, role-scoped, MFA enforced, logged, with access rights reviewed quarterly. International engagements have residency scoped per deployment.

SUBPROCESSOR	ROLE	SUBPROCESSOR	ROLE
Amazon Web Services	Infrastructure, storage	**Retell**	Voice conversation layer
Azure OpenAI	Model inference	**ElevenLabs, Cartesia**	Text to speech
Anthropic via Bedrock	Model inference	**Telnyx, Twilio**	Telephony
SendGrid	Transactional email		

Thirty days notice before we add or replace any subprocessor, with a right to object. Retention runs for the life of the engagement. On termination we provide a CSV export for thirty days, then delete everything including backups within ninety days.

3 Model training

No. Not your AR data, not prompts, not documents, not transactions, not outputs, not feedback. Not by AgentCollect and not by any provider in the chain. This is contractual on our side and flowed down to every provider, with **zero-retention inference endpoints** configured. No model anywhere is fine-tuned on client data. Base models are stateless at inference and receive only the context required for the conversation in progress, which is why no other client's data can enter yours and vice versa.

4 Governance and auditability

You define the rails: which actions the agent may take, minimum amounts, payment plan parameters, approval boundaries, and how disputes are handled against your own policies. Anything that falls outside those rails is never improvised, it routes to your team.

Every conversation is stored turn by turn as **text, audio, the reasoning trace, and the compliance checks applied before the agent spoke**. Any call or email can be reconstructed utterance by utterance. That is a structural property of the design: the pipeline is speech to text, then reasoning, then text to speech, rather than speech to speech, specifically so there is something to audit.

On human approval, the honest answer is that approval sits at the boundary rather than per action. If there is a specific action you want gated case by case, name it and we build that gate.

5 Failure and accountability

Compliance rules are checked on every turn **before** anything is sent, not after. Never obstructing a customer's route to a human is a hard design rule, not a setting. Escalated cases are reviewed by a licensed attorney, and we do not sue, garnish, or file in court.

Kill switches operate at three levels and stop new activity on the spot: a **single account**, a **campaign**, or the **entire program**. Engagements are month to month with no minimum, and one email ends it with a short transition window so no customer is left mid-conversation with a dead number.

On incident response, a suspected credential compromise or data exposure triggers containment first, then investigation, then written notice to your team with what we know and what we do not. Notification timelines run through the Data Processing Addendum.

0	0	0	168,000+
DATA BREACHES	CFPB COMPLAINTS	FDCPA VIOLATIONS	SEQUENCES RUN

CERTIFICATIONS

SOC 2 REPORT UNDER NDA

Microsoft SSPA SUPPLIER AUDIT PASSED

Amazon AVSP PENETRATION TEST PASSED

Regulation F / FDCPA ENFORCED AT SYSTEM LEVEL

END TO END: WHAT HAPPENS TO ONE INVOICE

01 Read
An open invoice, or a bank transaction, is read from your system into your tenant in AWS us-east-1.

02 Build context
Your tenant assembles the context for one conversation, inside your rails. Encrypted at rest, per-tenant.

03 Infer
That context goes to the model over a zero-retention endpoint and returns text. Nothing is retained provider side.

04 Reach out
Delivered by the disclosed voice, telephony or email provider for that channel, under your brand. Never ours.

05 Return
The outcome lands back in your tenant, and where you enable it, in your ERP. Payment goes directly to you.

RETAINED, IN YOUR TENANT
AR context, the conversation, its audio, its reasoning trace and compliance checks, for the life of the engagement. Deleted on your instruction. Voice, telephony and email providers hold delivery-level records under their own DPAs, which we will list for you.

NEVER
Retained at the inference layer, which runs on zero-retention endpoints. Used for training, by us or by any provider. Shared with another client. Sold. Nor does any funds movement pass through AgentCollect.

AVAILABLE UNDER NDA, SAME BUSINESS DAY

The **full SOC 2 Type II report**, the **penetration test executive summary**, the **subprocessor list with locations**, the **Data Processing Addendum**, and a completed security questionnaire. Request them at **agentcollect.com/security**. We supply the mutual NDA ourselves, so nothing waits on your legal team to draft one. Most enterprise procurement teams finish review in three to five business days.