ShopVoices

Security & Privacy Overview

A two-page brief for IT, security and data-protection reviewers · 2026

On-premise is the architecture, not a feature

ShopVoices is the in-store AI assistant that lets shoppers and staff ask for products in their own words and see the answer — price, stock and the route to the shelf — in seconds. Speech-to-text, the language model and the product search all run on the store's own hardware. No voice audio or transcript is ever sent to the cloud, and the spoken words are never stored.

How a shopper's voice is handled

1. **Capture, locally**
   The microphone feeds audio directly into the in-store system; on-device voice-activity detection notices when the shopper has finished.
2. **Transcribe, locally**
   Speech-to-text runs on the store's own hardware, in ten languages with automatic detection. Audio is held only in memory.
3. **Understand, locally**
   A language model on the same hardware works out what the shopper wants — no cloud call.
4. **Answer, locally**
   The system searches the store's own product, stock and layout database and shows the answer and route.
5. **Discard**
   The raw audio is processed in memory and discarded — never persisted, never transmitted. Neither is the transcript.

What is stored — and what never is

Data	How it is handled
Voice audio	Processed in memory on the store's own hardware. **Never persisted, never transmitted.**
The transcript (the words spoken)	**Never stored.** Used in memory to find the answer, then discarded.
Operational telemetry	Optional, anonymous and structurally an **allowlist**: timings, counts, category and intent tags, a language code, a success flag, a thumbs-up/-down. There is no field the spoken words could travel in. Permission-locked; can be switched off entirely.
Richer product-level analytics	A retailer-controlled configuration, never a default, used only where the applicable legal basis is in place — and it still carries no audio and no transcript.

> Because the guarantee is enforced by **where computation physically happens**, it can be verified in due diligence — on the store's own network, with your own tooling — rather than taken on trust from a policy document.

Layered security, no honeypot

- ✓ **Certificate-based device identity** — every kiosk, counter and handheld authenticates with mutual TLS; unknown devices cannot join.
- ✓ **Rotating, hashed credentials** — device tokens rotate and are stored only as hashes.
- ✓ **Signed software releases** — updates are cryptographically signed, rolled out in stages behind health gates, and roll back automatically.
- ✓ **Permission-locked local storage** — the small on-premise footprint is locked down at the file-system level.
- ✓ **Full audit trail** — administrative changes in the back office are logged and attributable.
- ✓ **No recordings to breach** — the most valuable target simply does not exist: there is no audio at rest anywhere in the system.

GDPR quick reference

The cloud voice-AI risk	Why it does not arise with ShopVoices
Cross-border transfer chain (Schrems II / SCCs)	**No audio transmission** — voice never leaves the building, so no audio transfer chain arises.
Article 28 processor & sub-processor agreements	**No cloud voice processor** — the answer path adds no processor chain for voice data.
Article 9 special-category biometrics	**No voiceprint** — ShopVoices transcribes a request to fulfil it; it does not identify people.
Retained recordings with incidental sensitive content	**Data minimisation by design** — audio is ephemeral and the transcript is never stored.
A large data footprint to govern	**A tiny, governable footprint** — anonymous metrics only, with an off switch.

EU AI Act orientation

The Act's prohibitions and its high-risk biometric-identification category target uses ShopVoices does not perform: it transcribes a request rather than recognising or profiling a person, performs no emotion recognition, and does not categorise people by sensitive traits. Analytics are never used to evaluate individual employees. As a local, task-specific assistant it is not a general-purpose-AI model provider.

Next step: request a technical & privacy briefing, or a focused pilot measured against your own baseline — **www.shopvoices.com/company/contact.html**. We are glad to support your DPIA and records-of-processing work with a small, clearly-described data footprint.