Vendor Risk Assessment — Pre-completed

Version: 2026-04 · **Owner:** Engineering · **Distributed as:** `Votriz_Vendor_Risk_Assessment.pdf`

This is the standard security questionnaire pre-filled. Drop it into your vendor management workflow as-is, or use it as a starting point for a custom Enterprise assessment via security@votriz.com.

Section A — Company

Question	Answer
Legal entity	Fortunato AI Inc.
Product	Votriz AI Presence OS
Headquarters	United States
Hosting region	US (EU available on Enterprise)
Year founded	2026
Number of employees handling customer data	1 (founder/CTO; designated AI Safety Officer)
Primary security contact	security@votriz.com

Section B — AI governance

Question	Answer
Does the AI follow established governance standards?	Aligned with NIST AI RMF (Govern / Map / Measure / Manage). Formal NIST + ISO 42001 attestation on the 2027 roadmap.
Documented model inventory?	Yes — `AI_MODEL_INVENTORY.md`.
Models developed in-house or third-party?	Third-party. Anthropic Claude Haiku (primary); OpenAI fallback; fal.ai for image generation. Votriz orchestrates, does not train.
Designated AI safety officer?	Yes — Founder/CTO.

Section C — Data handling

Question	Answer
Customer data used to train models?	**No.** Anthropic API contract excludes API data from training. Votriz does no fine-tuning.
Data segregation?	Three layers — application (`org_id` from JWT, per-query filter), database (PostgreSQL RLS policies on 9 tables), infrastructure (dedicated schema/instance available on Enterprise).
Proof of isolation?	`scripts/pre_prod_validation.sh` runs 15 multi-tenant assertions before every deploy; `GET /security/isolation-proof` returns live machine-readable evidence.
PII transmitted to LLMs?	Minimized. `services/pii_redactor.py` runs on chatbot + email-personalization paths. Lead generator extracts public business contact information by design (no redaction there).
Data retention policy?	Documented in `DATA_RETENTION_POLICY.md`. 30 days post-cancellation for account data; 7 years for audit logs (SOC 2). AI prompts ephemeral.
GDPR right of access?	`GET /settings/export-data` (Owner) — full JSON export of all org-scoped data.
GDPR right to erasure?	`POST /auth/deletion-request` queues a vetted human-reviewed deletion within 30 days.

Section D — Security controls

Question	Answer
Encryption in transit?	TLS 1.3 (Cloudflare edge).
Encryption at rest?	AES-256. OAuth tokens via per-row Fernet (key in hardware-backed secure storage); database via volume-level encryption.
Key management?	All keys held in hardware-backed secure storage and namespaced per service; never in code or configuration files.
RBAC?	5 built-in roles, 40+ permissions, brand-level scoping, custom roles on Enterprise.
MFA?	TOTP (RFC 6238) flow implemented this sprint. Mandatory enforcement available as Enterprise contract requirement.
Audit logging?	Append-only `security_audit_log` with database-trigger-enforced immutability. 7-year retention. Queryable + exportable via `/audit/log/*`.
Defenses against prompt injection?	`services/prompt_guard.py` pattern-matches common jailbreak markers. Output also passes through Brand DNA scoring + human approval gate.
Defenses against model extraction?	System prompts stored server-side; never sent to clients. Rate limiting on AI generation endpoints.
Penetration testing?	Internal automated suite + planned third-party engagement Q3 2026.

Section E — Operational resilience

Question	Answer
SLA target?	99.9% monthly uptime. Documented in `SLA.md`.
Status page?	votriz.com/status — 30-second polling.
Fallback mode if AI fails?	Multi-provider failover (Anthropic → OpenAI). Manual content / publishing / approval workflows function without AI.
Backup strategy?	Daily encrypted dumps + 30-day rolling retention + monthly restore drill.
Incident response plan?	Documented in `AI_INCIDENT_RESPONSE.md` — four severity levels, defined response times, kill switches at three layers.
Disaster recovery?	`scripts/disaster_recovery.sh` documents the full recovery procedure.

Section F — Intellectual property

Question	Answer
Who owns AI-generated content?	The customer. Votriz claims no IP rights over outputs generated through the platform.
Does Votriz claim rights to customer data?	No. Customer data is processed solely to deliver the service.
Copyright indemnity?	Outputs are original synthesis from brand-specific voice profiles, not reproduction of training data. Customer is responsible for review before publication; Votriz cooperates with copyright inquiries and removes flagged content promptly. Detailed terms in `votriz.com/legal/terms.html` §8.

Section G — Compliance & certifications

Question	Answer
GDPR compliance?	Active. DPA available. Sub-processor list published. Deletion + export endpoints live.
CAN-SPAM compliance?	Active. One-click unsubscribe on every campaign; sender identification on every send.
SOC 2?	Aligned with controls in place. Type I attestation Q4 2026. Type II Q2 2027. Readiness summary: `SOC2_READINESS.md`.
ISO 27001 / 42001?	On the 2027 roadmap.
HIPAA?	Not certified. Votriz is not currently a HIPAA business associate; do not store Protected Health Information.
PCI DSS?	Not directly — billing is handled by Stripe, who is PCI Level 1 certified. Votriz never sees raw card data.

Section H — Reporting concerns

Issue	Channel	Response time target
Security incident	security@votriz.com	1 hour (Enterprise), 4 hours (Agency) — see `SLA.md`
Privacy / DPA	privacy@votriz.com	5 business days
AI quality / bias	In-app "Report AI issue" button or `POST /ai/report-issue`	24 hours
Legal	legal@fortunatoai.com	5 business days

Section I — Custom requests

If your security team needs answers beyond what's covered here, or a tailored questionnaire format for your vendor management system, security@votriz.com is the right starting point. Enterprise customers receive

custom security reviews + direct access to the engineering team as part of the contract.