vvault

Security & Privacy Architecture

Technical Whitepaper — Vendor Security Review

ZERO CLOUD EXPOSURE
All processing occurs on your hardware. No external API calls after installation.

LOCAL AI MODEL
phi3:mini runs entirely on your machine. No data sent to OpenAI, Anthropic, or any third party.

GDPR BY ARCHITECTURE
Your security policies, questionnaire content, and answers never leave your network. Ever.

FULLY AUDITABLE
Complete audit trail. Every action logged with user, timestamp, and content.

Version 1.0

Published: April 2026

CONFIDENTIAL

For vendor security review purposes

getvvault.com | support@getvvault.com

// EXECUTIVE SUMMARY

What This Document Covers

This whitepaper provides a complete technical description of Vvault's security and privacy architecture. It is intended for CTOs, IT security teams, compliance officers, and procurement departments evaluating Vvault as a vendor tool for security questionnaire automation.

Vvault is a locally-deployed software application that automatically fills in SOC2 and security questionnaires using your company's own policy documents. The application runs entirely within your own infrastructure using Docker containers. After the initial installation, Vvault makes zero outbound network connections. Your security policies, questionnaire content, and generated answers never leave your network.

✔ **Core Privacy Guarantee**	Vvault processes all data locally. No content from your questionnaires, policy documents, or generated answers is transmitted to any external server, cloud service, or AI provider — including the developers of Vvault. This is verifiable using standard network monitoring tools.

Key Security Properties

Zero external data transmission. After installation, all processing is local. Network monitoring will show zero outbound connections during operation.

Local AI inference. The AI model (phi3:mini via Ollama) runs entirely on your hardware. No API calls are made to OpenAI, Anthropic, Microsoft, or any other AI provider.

Encrypted credentials. All passwords are hashed using bcrypt with salt. JWT tokens manage sessions. HTTPS enforced via self-signed certificate.

Role-based access control. Admin and Viewer roles with distinct permissions. Admins approve answers; Viewers can run autofill and review only.

Complete audit trail. Every approve, reject, and edit action is logged with username, timestamp, and full question/answer content.

Isolated containerisation. Each component runs in its own Docker container. The application, database, and AI model are fully isolated.

Data sovereignty. Your data is stored exclusively in a PostgreSQL database running on your own machine. You control it entirely.

// ARCHITECTURE

System Architecture Overview

Vvault is deployed as a set of Docker containers on your local machine or internal server. The architecture is intentionally designed to eliminate all external dependencies after the initial setup phase.

Component Architecture

Component	Technology	Purpose	Network Access
Frontend	Nginx + React	Web dashboard served locally on port 3443 (HTTPS)	None — local only
Backend API	FastAPI (Python)	REST API handling all business logic and orchestration	None — local only
AI Model Runner	Ollama	Runs phi3:mini language model for answer generation	None after install
Embedding Model	Ollama (nomic-embed-text)	Generates vector embeddings for semantic search	None after install
Database	PostgreSQL with pgvector	Stores answers, knowledge base, users, and audit logs	None — local only
Reverse Proxy	Nginx	HTTPS termination and request routing	None — local only

Network Architecture

The following describes every network connection Vvault makes, categorised by phase:

Installation Phase (one-time only)

Destination	Domain	Purpose	Frequency
Docker Hub	hub.docker.com	Pull base images for containers	One-time
Ollama Registry	ollama.com	Download phi3:mini model (~2.3 GB)	One-time
Ollama Registry	ollama.com	Download nomic-embed-text model (~270 MB)	One-time
PyPI	pypi.org	Install Python dependencies during build	One-time

Runtime Phase (ongoing operation)

After installation is complete, Vvault makes ZERO outbound network connections. All runtime communication is internal between Docker containers on localhost.

✔ Runtime Network Verification	You can verify this independently using Wireshark, Little Snitch (macOS), or Windows Firewall logs. Filter for the vvault_backend and vvault_ollama containers. You will observe zero outbound connections to any external IP address during normal operation.

// DATA FLOW

How Your Data Is Processed

This section documents exactly what happens to your data at each stage of the Vvault workflow. Understanding this flow is essential for GDPR compliance assessment and internal security reviews.

Step 1

Policy Document Upload

You upload your security policy documents (PDF, TXT, or DOCX) via the web dashboard. Files are received by the FastAPI backend running locally. Documents are chunked into semantic segments and stored in the local PostgreSQL database. The nomic-embed-text model generates vector embeddings for each chunk, also stored locally. Original files are processed in memory and not written to disk beyond the database. No content leaves the machine at any point.

Step 2

Questionnaire Upload

You upload an Excel, Word, or PDF questionnaire. The backend parses the file to extract question rows. Questions are identified using column header detection and content-based heuristics. All parsing occurs in memory on your machine. The questionnaire file itself is not stored permanently — only the extracted questions are retained in the session.

Step 3

Answer Generation

For each question, Vvault follows a three-tier lookup: (1) Semantic cache lookup — searches previously approved answers using vector similarity in the local database. (2) Template matching — matches against 20 pre-defined SOC2 answer templates using keyword and semantic matching. (3) LLM generation — if no cache or template match, the local phi3:mini model generates an answer using retrieved policy context. All three tiers operate entirely locally. The LLM receives only the question and relevant policy excerpts as context — no personal data, credentials, or sensitive metadata is included in the prompt.

Step 4

Human Review

All generated answers are presented in the review dashboard with confidence scores and source attribution. No answers are automatically submitted to anyone. Your team approves, edits, or rejects each answer. Every action is logged in the audit trail.

Step 5

Export

The completed questionnaire is written back to an Excel file and downloaded by your browser. The file never passes through any external server. The download is a direct browser download from your locally-running application.

// AI MODEL

Local AI Model — Technical Details

Vvault uses two AI models, both running locally via Ollama. Neither model makes external API calls during operation. This section provides technical details for security review.

Language Model — phi3:mini

Property	Detail
Model name	phi3:mini (3.8B parameters)
Developer	Microsoft Research
License	MIT License — permits commercial use
Model size	Approximately 2.3 GB on disk
Runtime	Ollama — open source local model runner
Hardware requirement	8 GB RAM minimum, 16 GB recommended
External calls	Zero — model runs entirely on local CPU/GPU
Training data	Publicly available text datasets — no customer data in training
Data retention	Model receives prompt at inference time only — no conversation history persisted in model
Inference location	http://ollama:11434 — internal Docker network only

Embedding Model — nomic-embed-text

Property	Detail
Model name	nomic-embed-text v1.5
Developer	Nomic AI
License	Apache 2.0 — permits commercial use
Model size	Approximately 270 MB on disk

Purpose	Generates vector embeddings for semantic similarity search
External calls	Zero — runs entirely locally via Ollama
Output	768-dimensional floating point vectors stored in PostgreSQL

What the AI Models See

The language model (phi3:mini) receives the following information as input during answer generation:

- The security questionnaire question text
- Relevant excerpts from your uploaded policy documents (maximum 2,000 characters)
- A system instruction defining the SOC2 expert role

The AI models do NOT receive:

- Your company name, employee names, or organisational details
- Customer names or client-sensitive information
- Database credentials, API keys, or authentication tokens
- Historical answers from other companies
- Any data from other Vvault installations

■ Important Limitation

Like all language models, phi3:mini can generate inaccurate answers. Vvault flags low-confidence answers for mandatory human review. All answers require approval before use. Vvault is a drafting assistant — final responsibility for answer accuracy remains with your compliance team.

// DATA STORAGE

Data Storage & Retention

All Vvault data is stored in a PostgreSQL database running in a Docker container on your machine. This section documents every data category stored and how it is handled.

Database Schema — What Is Stored

Table	Data Stored	Retention	Encrypted
qa_cache	Questions, generated answers, confidence scores, source attribution, audit metadata, vector embeddings	Until manually deleted or docker-compose down -v	At rest via filesystem
knowledge_base	Chunked text from your policy documents, vector embeddings, source filename	Until manually deleted	At rest via filesystem
users	Username, bcrypt password hash, role (admin/viewer)	Until user deleted	Passwords hashed — never stored plaintext
audit_logs	Username, action (approve/reject), question text, answer text, timestamp, run ID	Permanent until reset	At rest via filesystem

Password Storage

Passwords are never stored in plaintext. Vvault uses bcrypt with a randomly generated salt for each password. The bcrypt work factor makes brute-force attacks computationally expensive. Password hashes are stored in the users table and are never transmitted or logged.

Session Management

Sessions are managed using JSON Web Tokens (JWT) with HMAC-SHA256 signing. Tokens expire after 12 hours. The JWT secret is configured by the administrator in the .env file during setup. Tokens are transmitted over HTTPS only. No session data is stored server-side — token validation is stateless.

Vector Embeddings

Policy document chunks and question texts are converted to 768-dimensional vector embeddings using nomic-embed-text. These vectors are stored in PostgreSQL using the pgvector extension. Embeddings

enable semantic similarity search without requiring exact keyword matching. Embeddings are mathematical representations — they cannot be directly reversed to reconstruct the original text with certainty, though they should be treated as sensitive data.

Complete Data Removal

To completely remove all Vvault data from your system:

1. Stop all containers: docker-compose down
2. Remove all data volumes: docker-compose down -v
3. Remove Docker images: docker rmi $(docker images | grep vvault | awk '{print $3}')
4. Remove the installation directory: rm -rf /path/to/vvault
5. Verify removal: docker volume ls (no vvault volumes should appear)

// ACCESS CONTROL

Authentication & Authorisation

User Roles

Role	Permissions	Restrictions
Admin	Run autofill, review answers, approve answers, reject answers, manage users, create users, delete users, download completed questionnaires, view audit logs	None
Viewer	Run autofill, review answers, download completed questionnaires	Cannot approve or reject answers. Cannot manage users. Cannot view audit logs.

Authentication Flow

Credential submission. Username and password submitted via HTTPS POST to /auth/login

Password verification. bcrypt.checkpw() compares submitted password against stored hash

Token issuance. On success, JWT token issued containing username, role, and 12-hour expiry

Request authentication. All subsequent API requests require Authorization: Bearer header

Token validation. Backend middleware validates JWT signature and expiry on every request

Role enforcement. Route handlers check request.state.role before executing privileged operations

Default Credentials

■ Action Required on First Login

The default admin password is 'changeme123'. Vvault detects this on first login and forces an immediate password change before the application opens. The default password must be changed before the system can be used. Do not deploy Vvault in a production environment without completing this step.

Network Access Controls

Vvault is accessible only on localhost by default. The Nginx reverse proxy listens on port 3443 (HTTPS). No ports are exposed to the public internet unless explicitly configured in docker-compose.yml. For multi-user deployments on a local network, appropriate firewall rules should be configured to limit access to authorised workstations only.

// COMPLIANCE

GDPR & Regulatory Compliance

Vvault's architecture is designed to support GDPR compliance by eliminating external data transmission entirely. This section addresses the key GDPR considerations relevant to Vvault deployment.

Data minimisation (Art. 5(1)(c))

Vvault stores only what is necessary for operation. Policy document chunks, questions, answers, user credentials, and audit logs. No telemetry, usage analytics, or behavioural data is collected.

Data localisation

All data remains within your infrastructure. Vvault does not transfer data to third countries. The application has no mechanism to transmit data externally after installation.

Right to erasure (Art. 17)

Complete data removal is achievable via docker-compose down -v. This permanently destroys all data including the database, knowledge base, and answer library. This operation is irreversible.

Data controller status

Your organisation is the data controller for all content processed by Vvault. Vvault (the software) is a data processor tool that operates entirely under your control. The developers of Vvault have no access to your data.

Sub-processor assessment

Vvault has no sub-processors. The AI models (phi3:mini and nomic-embed-text) run locally and are not provided as a service by Microsoft or Nomic AI during operation — they are downloaded once and run locally under your control.

Data Processing Agreement

Because Vvault processes no data on behalf of the vendor (all processing is local), a traditional DPA between your organisation and Vvault may not be required. Consult your legal counsel for a definitive assessment based on your specific deployment.

Relevant Frameworks

Framework	Relevance	Vvault Position
GDPR	European data protection regulation	Supported by architecture — zero external transmission
SOC2 Type II	Security, availability, confidentiality controls	Vvault assists with questionnaire completion — not itself SOC2 certified

ISO 27001	Information security management	Local deployment supports data confidentiality and integrity requirements
NIS2 Directive	EU cybersecurity directive for essential/important entities	Local-only architecture reduces third-party risk exposure
UK GDPR	Post-Brexit UK data protection	Same architectural compliance as EU GDPR

// VERIFICATION

How to Verify Zero Data Transmission

Vvault's core privacy claim — that no data leaves your machine during operation — is independently verifiable. This section provides specific instructions for your IT security team to verify this claim using standard network monitoring tools.

Method 1 — Wireshark (All platforms)

1. Install Wireshark from wireshark.org
2. Start a capture on your primary network interface
3. Apply filter: ip.addr != 127.0.0.1 && ip.addr != 172.16.0.0/12 (excludes localhost and Docker internal network)
4. Use Vvault normally — upload a policy document, run autofill on a questionnaire
5. Observe captured packets — you should see zero outbound connections to external IP addresses
6. The only traffic during operation will be between your browser and localhost:3443

Method 2 — Docker Network Inspection

Run the following commands to inspect container network activity:

List all container network connections:

```
docker stats --no-stream
```

Inspect outbound connections from the backend container:

```
docker exec vvault_backend netstat -tn | grep ESTABLISHED
```

Expected result: connections will show only internal Docker network addresses (172.x.x.x range) and localhost. No external IP addresses should appear.

Method 3 — macOS Little Snitch

Install Little Snitch and create a rule to monitor all connections from Docker Desktop. During Vvault operation you should observe zero connection attempts to external addresses. Any connection attempt will be shown in the Little Snitch Network Monitor.

Method 4 — Windows Firewall Logging

1. Enable Windows Firewall logging: Start > Windows Defender Firewall > Advanced Settings > Properties > Logging
2. Set log dropped packets and successful connections to Yes
3. Use Vvault for 10 minutes

4. Review log at C:\Windows\System32\LogFiles\Firewall\pfirewall.log

5. Filter for Docker Desktop process — outbound connections should show only localhost and internal Docker subnet

✔ **Verification Support**	If your IT security team requires assistance setting up network monitoring or interpreting results, contact support@getvvault.com. We offer a free 30-minute screen share session to walk through verification with your team.

// SECURITY PRACTICES

Secure Development & Operational Security

Dependency Management

Vvault uses the following key dependencies. All dependencies are open source with established security track records:

Dependency	Version	Purpose	License
FastAPI	Latest stable	Python web framework	MIT
PostgreSQL	15+	Primary database	PostgreSQL License
pgvector	Latest stable	Vector similarity search extension	PostgreSQL License
Ollama	Latest stable	Local AI model runtime	MIT
phi3:mini	3.8B	Answer generation language model	MIT
nomic-embed-text	v1.5	Text embedding model	Apache 2.0
bcrypt	Latest stable	Password hashing	Apache 2.0
python-jose	Latest stable	JWT token handling	MIT
Nginx	Latest stable	Reverse proxy and HTTPS	BSD
Docker	20.10+	Container runtime	Apache 2.0

Known Limitations & Accepted Risks

Self-signed TLS certificate

Vvault uses a self-signed certificate for HTTPS. Browsers will show a security warning on first access. This is expected and does not reduce security for localhost deployments. Enterprise customers can replace with a certificate from their internal CA.

No automatic updates

Vvault does not auto-update. Updates require manual download and docker-compose up --build. This is intentional — automatic updates would require external connections.

AI answer accuracy

phi3:mini is a small language model. Answer quality depends on the quality and completeness of uploaded policy documents. Human review of all answers is mandatory and enforced by the workflow.

Single-node deployment

Vvault v1.0 runs on a single machine. High availability and clustering are not supported. This is appropriate for the intended use case of periodic questionnaire completion.

// INCIDENT RESPONSE

Security Incident Reporting & Response

Vvault is committed to responsible security disclosure and prompt response to reported vulnerabilities.

Reporting a Security Vulnerability

If you discover a security vulnerability in Vvault, please report it responsibly:

- Email: support@getvvault.com with subject line 'Security Vulnerability Report'
- Include: Description of the vulnerability, steps to reproduce, potential impact assessment
- Do not: Publicly disclose the vulnerability before we have had the opportunity to address it
- Response commitment: We will acknowledge receipt within 24 business hours
- Resolution commitment: We will provide a fix or mitigation within 30 days for critical issues

Local Deployment Security Responsibilities

Because Vvault runs on your own infrastructure, certain security responsibilities fall to your organisation as the operator:

Operating system security. Keep the host operating system patched and updated. Vvault containers do not manage host OS security.

Network access control. Restrict access to port 3443 using firewall rules. Only authorised users on your network should be able to reach the Vvault dashboard.

Backup. Back up the PostgreSQL database regularly using the provided backup scripts. Vvault does not perform automated backups.

Admin password. Change the default admin password immediately on first login. Use a strong, unique password stored in your password manager.

Docker security. Follow Docker security best practices for your deployment environment. Ensure Docker Desktop is kept updated.

// INSTALLATION

Secure Installation Guide

This section provides security-focused guidance for IT administrators deploying Vvault.

Pre-Installation Checklist

- ■ Host machine meets minimum requirements: 8 GB RAM, 10 GB free disk, Docker Desktop 4.0+
- ■ Docker Desktop is running and up to date
- ■ Port 3443 is not in use by another application
- ■ You have generated a secure JWT_SECRET (minimum 32 characters, random)
- ■ You have set a strong DB_PASSWORD in the .env file
- ■ Internet connection available for initial Docker image and model download

Environment Configuration Security

The .env file contains sensitive configuration. Secure it appropriately:

• DB_PASSWORD: Set to a strong random password. Minimum 16 characters. Never use the default.

• JWT_SECRET: Generate using: python -c "import secrets; print(secrets.token_hex(32))" — minimum 32 characters

• File permissions: Restrict .env file permissions to the running user only (chmod 600 .env on Linux/macOS)

• Version control: Never commit .env to Git. The .gitignore file excludes it by default.

Post-Installation Security Steps

1. Log in and change the default admin password immediately
2. Create individual user accounts for each team member — do not share the admin account
3. Test network isolation using Method 1 or 2 from the Verification section
4. Configure host firewall to restrict port 3443 access to authorised workstations
5. Schedule regular database backups using the provided backup scripts
6. Note your installation date — plan to review and update Vvault quarterly

// FAQ

Security Review — Frequently Asked Questions

Does Vvault send any data to your servers?

No. After installation, Vvault makes zero outbound network connections. The developers of Vvault have no access to your data, your questions, your policy documents, or your answers. This is verifiable using the network monitoring methods described in this document.

Is the AI model provided as a cloud service?

No. The AI models (phi3:mini and nomic-embed-text) are downloaded once during installation and run entirely on your hardware using Ollama. There are no API calls to Microsoft, Nomic AI, or any other provider during operation.

Who can access our data?

Only users with valid Vvault credentials on your local network. The admin account controls user creation. Vvault developers, support staff, or any external party have zero access to your Vvault data.

What happens to our data if we cancel or stop using Vvault?

Your data remains on your machine until you choose to delete it. Running docker-compose down -v permanently and irreversibly removes all Vvault data. You are always in control.

Is Vvault SOC2 certified?

Vvault v1.0 is not SOC2 certified. As an early-stage product, certification is on the roadmap. However, because Vvault is a locally-deployed tool with zero external data transmission, many SOC2 concerns that apply to cloud tools are architecturally eliminated.

Can Vvault be deployed on an air-gapped network?

Yes. After the initial installation (which requires internet for Docker image and model download), Vvault operates with zero internet connectivity. It is suitable for air-gapped environments.

Does Vvault support SSO or Active Directory integration?

Not in v1.0. User management is handled within Vvault using local credentials. SSO/LDAP integration is on the product roadmap.

What is the data residency of our policy documents?

Your policy documents are stored exclusively on the machine where Vvault is installed. They are never transmitted to any external location. Data residency is entirely determined by where you choose to run the Docker containers.

Can we review the source code?

The Vvault codebase is available for review under NDA for enterprise customers conducting security assessments. Contact support@getvvault.com to arrange code review access.

What support do you provide for security incidents?

Report security vulnerabilities to support@getvvault.com. We commit to acknowledging reports within 24 business hours and providing critical fixes within 30 days.

// CONTACT

Contact & Further Information

For questions not addressed in this whitepaper, or to arrange a technical security review call with our team, please use the following contacts:

Enquiry Type	Contact	Response Time
General security questions	support@getvvault.com	Within 24 business hours
Vulnerability reports	support@getvvault.com	Within 24 business hours
Enterprise security review / screen share	girish@getvvault.com	Book at getvvault.com/demo
Source code review requests	girish@getvvault.com	Within 48 business hours
Procurement / vendor assessment support	girish@getvvault.com	Within 24 business hours

✔ Free Security Review Call

We offer a free 30-minute screen share session for IT security teams evaluating Vvault. We will walk through the architecture, demonstrate network isolation verification, and answer any technical questions your team has. Book at getvvault.com/demo.

Vvault Security & Privacy Architecture — Technical Whitepaper v1.0

Published April 2026. getvvault.com | support@getvvault.com

This document is provided for vendor security evaluation purposes. Information is accurate as of the publication date and subject to change as the product evolves. For the latest version of this document visit getvvault.com/security.