FIMIL, INC.

VSA Full v2021

Vendor Security Alliance Full questionnaire — comprehensive security assessment covering data protection, policies, proactive and reactive security, software supply chain, and compliance.

March 20, 2026

Security Questionnaire Response

trust.fimil.dev

security@fimil.dev

Response Summary

Answer	Count	Percentage
Yes	83	74.1%
Partial	17	15.2%
No	11	9.8%
N/A	1	0.9%
Total	112	100%

Service Overview

8/12 controls satisfied

ID	Question	Answer	Explanation
SO-1	Company name.	YES	Fimil, Inc. — a Delaware C Corporation.
SO-2	Describe the service being provided.	YES	Fimil is a unified application security platform that orchestrates open-source security scanners (SAST, SCA, Secrets, IaC, Container) behind a single dashboard with finding deduplication, priority scoring, and policy enforcement.
SO-3	What technology stack is used to provide the service?	YES	Python/FastAPI backend, React/TypeScript frontend, PostgreSQL database, Redis for caching and job queuing, Celery workers, Docker containers for scanner isolation, Kubernetes on DigitalOcean for orchestration.
SO-4	Is the service hosted in your own datacenter, in the cloud, or deployed on-premise at the customer's location?	YES	Cloud-hosted on DigitalOcean Kubernetes (SaaS). An Enterprise self-hosted deployment model is also available for on-premise installation.
SO-4a	If hosted in a data center, list all data center locations.	YES	DigitalOcean managed Kubernetes cluster; primary region is a US-based DigitalOcean data center. Single-region deployment currently.
SO-4b	What cloud providers do you rely on?	YES	DigitalOcean (infrastructure, managed PostgreSQL, container registry), Cloudflare (CDN/DDoS protection), Stripe (payment processing), Resend (transactional email), PostHog (analytics).
SO-4c	Have you researched the security best practices of your cloud provider(s)?	YES	Vendor Risk Management Policy (FIMIL-VRM-001) classifies all vendors into three tiers with documented security assessments. DigitalOcean is Tier 1 with SOC 2 report collection tracked.
SO-5a	Do you have the most recent penetration test report available?	NO	No external penetration test has been conducted. This is the primary remaining gap identified in our ISO 27001 assessment (A.5.35 FAIL).
SO-5b	Does the penetration test follow an industry approved methodology (e.g., OWASP, PTES)?	NO	No penetration test has been performed yet. When conducted, it will follow OWASP methodology.
SO-5c	Do you have Information Security Policies and Procedures available for review?	YES	Comprehensive policy suite available: ISMS Policy (FIMIL-ISMS-001), Access Control Policy, Data Governance Policy, Change Management Policy, Incident Response Plan, People Security Policy, and Vendor Risk Management Policy.
SO-5d	Do you have a Data Flow Diagram available for review?	PARTIAL	Architecture and data flow documented in internal technical documentation covering scanner pipeline and storage layers. No standalone visual data flow diagram maintained yet.
SO-5e	Do you have PCI, SOC 2 Type II, or ISO 27001 certification reports available for review?	NO	ISO 27001 and SOC 2 controls are implemented with a comprehensive policy suite, but certification audits have not yet been completed. Certification is planned.

Data Protection & Access Controls

20/23 controls satisfied

ID	Question	Answer	Explanation
DPAC-1	What customer data is required to provide the service (personal, financial, confidential, sensitive, government)?	YES	Fimil collects user email and name (personal), OAuth tokens for Git provider integration (confidential), and temporarily processes source code for scanning (ephemeral, never persisted). No financial, sensitive, or government data is collected.
DPAC-2	Do you have a data classification matrix available?	YES	Data Governance Policy (FIMIL-DGP-001) establishes a four-level classification: Public, Internal, Confidential (user PII, scan results), and Restricted (encryption keys, OAuth tokens, API credentials).
DPAC-3	How is customer data encrypted?	YES	TLS 1.2+ for data in transit. MultiFernet encryption (AES-128-CBC + HMAC-SHA256) with versioned key rotation for sensitive fields at rest. Database encryption via DigitalOcean managed PostgreSQL. Backups encrypted in S3.
DPAC-4	How does your organization decide who has access to sensitive data?	YES	Five-level RBAC hierarchy (Operator > Admin > Security > Developer > Viewer) enforced at every API endpoint via FastAPI dependencies. Tenant isolation ensures customers only access their own data.
DPAC-5	Do you have capabilities to anonymize data?	YES	Source code is ephemeral (clone-scan-delete, never persisted). Scan findings are normalized through the scanner pipeline, stripping source context. Secret redaction removes sensitive values before storage.
DPAC-6	How is anonymized data used?	YES	Anonymized scan findings are used solely for vulnerability reporting and trend analysis within the customer's own tenant. No cross-tenant data sharing or aggregation.
DPAC-7	What are the general rules for role provisioning, deprovisioning, and recertification?	YES	Access Control Policy governs provisioning. Email verification required for registration. Admin approval required for role elevation. People Security Policy includes offboarding checklists with access revocation. User deactivation immediately blocks all authentication.
DPAC-8	Which staff groups have access to personal or sensitive data?	YES	Currently sole founder with full system access. RBAC enforces that only Admin and Operator roles can access user management. Security role can access findings. Developer and Viewer roles have limited access.
DPAC-9	Is any sensitive data kept in hard copy?	NO	Fimil is a fully cloud-based platform. No hard copy sensitive data exists. All data is stored digitally with encryption.
DPAC-10	Is there a procedure in place for securely destroying hard copy sensitive data?	NA	No hard copy sensitive data exists. Fimil is entirely cloud-hosted with no physical data storage.
DPAC-11	Do you support secure deletion (degaussing/cryptographic wiping) of archived or backed-up data?	YES	Data Governance Policy documents data disposal procedures by classification level. Source code is ephemeral. Report retention configurable at 30 days. API tokens revocable. Session data expires automatically.
DPAC-12	Under what circumstances is customer data allowed to leave production systems?	YES	Customer data may leave production only via encrypted backups to S3, authenticated API exports by authorized users, and DSAR fulfillment. Source code never persists beyond the scan lifecycle.
DPAC-13	Do you have an internal password policy?	YES	Password policy enforces minimum 12 characters with complexity requirements (mixed case, digit, special character). Account lockout after 5 failed attempts with 30-minute cooldown.
DPAC-14	Are there complexity or length requirements for passwords?	YES	Passwords must be at least 12 characters with mixed case, at least one digit, and at least one special character. Maximum length allows 64+ characters.
DPAC-15	How are passwords hashed?	YES	Passwords are hashed using Argon2id (memory-hard) with per-password salt. Migrated from bcrypt to Argon2id for stronger resistance against GPU and ASIC attacks.
DPAC-16	Is a remote connection to the production systems (VPN) required?	YES	Production systems run on DigitalOcean Kubernetes with access restricted via Kubernetes RBAC and kubectl authentication. No direct SSH access; all management through authenticated Kubernetes API.
DPAC-17	Is MFA required for employees/contractors to log in to production systems?	YES	TOTP-based MFA with recovery codes implemented in the application with two-step login flow and encrypted secret storage. Production Kubernetes access controlled via DigitalOcean authentication and Kubernetes RBAC.
DPAC-18	Do your internal applications leverage SSO?	YES	OAuth2/OIDC federation supported with GitHub and generic OIDC providers for user authentication. SSO available on all plans.
DPAC-19	Which processors (vendors) have access to customer information?	YES	Sub-processor list published in trust center: DigitalOcean (infrastructure), Stripe (payments), Resend (email), PostHog (analytics). DPAs executed with all vendors.

ID	Question	Answer	Explanation
DPAC-20	Do your processors comply with your security standards?	YES	Vendor Risk Management Policy (FIMIL-VRM-001) defines three-tier classification with security requirements. All Tier 1 vendors (DigitalOcean, GitHub, Stripe) maintain SOC 2 certification.
DPAC-21	How do you regularly audit your critical vendors?	YES	Vendor Risk Management Policy defines annual review for Tier 1 vendors and biannual for Tier 2. SOC 2 report collection tracked. Exit strategies documented for critical vendors.
DPAC-22	Do you process EU citizens' personal data?	YES	Yes, Fimil is available globally and may process EU citizens' data. GDPR compliance requirements are tracked in the Compliance Register, and a Data Processing Agreement is published at /legal/dpa.
DPAC-23	Have you appointed a Data Protection Officer?	PARTIAL	ISMS Policy designates the CEO (sole founder) with data protection responsibilities. A formal DPO appointment will be formalized as the organization scales.

Policies & Standards

8/12 controls satisfied

ID	Question	Answer	Explanation
PS-1	Do you have a formal Information Security Program in place?	YES	ISMS Policy (FIMIL-ISMS-001) establishes the formal Information Security Management System aligned with ISO 27001:2022, supported by a comprehensive policy suite.
PS-2	Do you review your InfoSec policies at least annually?	YES	ISMS Policy defines annual policy review cadence. Statement of Applicability (FIMIL-SOA-001) tracks control implementation status. Policies are version-controlled in Git.
PS-3	Do you have an information security risk management program?	YES	Risk Assessment (FIMIL-RISK-001) uses a 5x5 likelihood-impact matrix identifying 15 risks with documented treatment plans and timelines (Q2-Q4 2026).
PS-4	Is there management support and/or a security management forum?	PARTIAL	ISMS Policy establishes management commitment to information security. Currently sole founder serving all governance roles; security steering committee to be established as team scales.
PS-5	Do you have a dedicated information security team?	PARTIAL	Currently sole founder with deep security domain expertise (building an application security platform). Dedicated security team to be established as the organization grows.
PS-6	Do your InfoSec and privacy policies align with industry standards (e.g., ISO 27001, NIST, SOC 2)?	YES	Policies are aligned with ISO 27001:2022 and SOC 2 Type II Trust Services Criteria. Compliance Register tracks regulatory obligations including GDPR and CCPA.
PS-7	Do you have a policy exception process?	YES	Change Management Policy (FIMIL-CHG-001) includes exception handling for emergency changes. ISMS Policy provides a framework for policy exceptions with documented risk acceptance.
PS-8	Do you have a formal disciplinary policy for security policy violations?	YES	People Security Policy (FIMIL-PPL-001) defines a disciplinary process for security policy violations. Currently sole founder; enforcement framework ready for team scaling.
PS-9	Do you perform background verification for all employment candidates?	PARTIAL	People Security Policy documents background verification requirements for all roles including security competence criteria. Not yet exercised as currently sole founder.
PS-10	Are all personnel required to sign Confidentiality Agreements?	PARTIAL	People Security Policy mandates confidentiality agreements for all personnel. Policy documented and ready for execution; not yet exercised as currently sole founder.
PS-11	Are all personnel required to sign an Acceptable Use Policy?	YES	Acceptable Use Policy published at /legal/acceptable-use. People Security Policy requires all personnel to acknowledge and sign the AUP.
PS-12	Are there procedures for change in employment or termination, including access revocation?	YES	People Security Policy includes offboarding checklists with access revocation procedures. Technical controls support immediate user deactivation, session invalidation, and API token revocation.

Proactive Security

19/29 controls satisfied

ID	Question	Answer	Explanation
PRO-1	How is network security testing performed (internal, third party, cadence)?	PARTIAL	Network security relies on Kubernetes network policies and scanner container isolation (--network=none). No independent network penetration testing has been performed yet.
PRO-2	How is application security testing performed?	YES	CI pipeline runs SAST (Semgrep, Bandit), container scanning (Trivy), linting, and automated tests on every push/PR. Fimil scans its own repositories through the platform. Security-specific tests in api/tests/security/.
PRO-3	What are your network vulnerability management processes?	PARTIAL	Container image scanning with Trivy in the deployment pipeline. Kubernetes network policies restrict traffic. No dedicated network vulnerability scanning tool deployed yet.
PRO-4	What is the timeframe for patching critical vulnerabilities?	YES	Formal patch management SLA: Critical 24h, High 7d, Medium 30d. Trivy blocks deployment of containers with critical vulnerabilities. EPSS enrichment prioritizes actively exploited CVEs. Dependabot provides automated dependency update PRs.
PRO-5	What tools do you use for vulnerability management?	YES	Fimil's own platform orchestrates Semgrep, Bandit, Trivy, Grype, OSV-Scanner, Gitleaks, TruffleHog, Checkov, Hadolint, and Syft for comprehensive vulnerability detection.
PRO-6	What are your application vulnerability management processes?	YES	Automated scanning in CI/CD with priority scoring (0-100) combining severity, age, reachability, and EPSS. Auto-triage rules classify findings. Finding deduplication groups equivalent findings across tools.
PRO-7	What tools do you use for application vulnerability management?	YES	Fimil platform with 12 integrated scanners covering SAST, SCA, Secrets, IaC, Container, and SBOM. EPSS enrichment and reachability analysis for prioritization.
PRO-8	How regularly do you evaluate patches and updates?	YES	Dependabot configured across all repos for automated dependency updates. Container image scanning on every build detects known CVEs. Patch management SLA: Critical 24h, High 7d, Medium 30d.
PRO-9	Do you have a responsible disclosure path published?	YES	SECURITY.md published with responsible disclosure policy, testing scope, safe harbor provisions, and security@fimil.dev contact with 48-hour acknowledgment SLA.
PRO-10	Do you have an established bug bounty program?	NO	No formal bug bounty program. Responsible disclosure policy is published via SECURITY.md but does not include financial incentives.
PRO-11	Are endpoint laptops centrally managed?	PARTIAL	Currently sole founder; no formal MDM solution deployed. Endpoint security policies are documented in the People Security Policy for future team scaling.
PRO-12	Do you have a standard device security configuration?	PARTIAL	People Security Policy documents device security requirements including full-disk encryption and screen lock. Formal standard configuration to be enforced via MDM as team grows.
PRO-13	Is sensitive or private data stored on endpoint devices?	NO	Customer data resides in cloud infrastructure (DigitalOcean managed PostgreSQL). Source code is ephemeral. No sensitive data is stored on endpoint devices.
PRO-14	How do you limit data exfiltration from production endpoints?	YES	Production access is via Kubernetes RBAC only; no direct SSH. Scanner containers run with --network=none. Application-level intrusion detection includes bulk export detection and API anomaly detection.
PRO-15	Do you have systems to mitigate web application vulnerabilities (WAF, proxies)?	YES	Cloudflare WAF deployed with managed rulesets for web application protection. Cloudflare provides DDoS protection. Application-level protections include rate limiting, CSRF protection, input validation, and Content-Security-Policy headers.
PRO-16	Do you have breach detection or anomaly detection with alerting?	YES	Security alert system provides automated threat detection: brute force detection, credential stuffing detection, API anomaly detection, and suspicious scan pattern detection. Falco provides runtime container integrity monitoring.
PRO-17	Are hosts uniformly configured?	YES	All production workloads run as immutable Docker containers on Kubernetes with standardized configurations: non-root, read-only filesystem, cap_drop ALL, no-new-privileges, defined resource limits.
PRO-18	Are production changes reviewed by at least two engineers?	PARTIAL	Branch protection enforced: GPG-signed commits required, CI status checks required, enforce admins enabled. CODEOWNERS file established. CI pipeline gates on linting, tests, and container scanning. Currently sole founder; multi-reviewer approval to be implemented as team grows.
PRO-19	What is your secrets management strategy?	YES	Kubernetes Sealed Secrets encrypt production credentials. MultiFernet encryption with versioned key rotation for sensitive database fields. Secret

ID	Question	Answer	Explanation
			redaction in logs. Gitleaks and TruffleHog scan for leaked secrets in CI.
PRO-20	Are all security events in production logged?	YES	Audit logging tracks 40+ action types with user_id, IP, timestamp, action, resource_type, and resource_id. Structlog JSON output with request correlation IDs. Impersonation tracked with both admin and impersonated user IDs.
PRO-21	Is the production network segmented?	YES	Kubernetes network policies segment the production network. Scanner containers run with --network=none for complete isolation. Ingress-only access with TLS termination.
PRO-22	Is there a process for network configuration changes?	YES	Change Management Policy (FIMIL-CHG-001) governs all infrastructure changes including network configuration. Helm-based deployments with atomic rollback ensure safe changes.
PRO-23	Is network traffic over public networks encrypted?	YES	TLS 1.2+ enforced for all public-facing traffic via cert-manager with Let's Encrypt. HSTS headers enabled. HTTP redirected to HTTPS.
PRO-24	What cryptographic frameworks are used for data in transit?	YES	TLS 1.2+ managed by cert-manager with Let's Encrypt certificates. HSTS enforced. Auth cookies set with Secure flag.
PRO-25	What cryptographic frameworks are used for data at rest?	YES	MultiFernet encryption (AES-128-CBC + HMAC-SHA256) with versioned key rotation for sensitive fields (OAuth tokens, API credentials). Database encryption via DigitalOcean managed PostgreSQL. S3 backups encrypted.
PRO-26	What cryptographic frameworks are used for passwords?	YES	Argon2id (memory-hard) with per-password salt for user passwords. SHA-256 hashing for API tokens and email verification tokens. Constant-time comparison for all secret operations.
PRO-27	Are any custom cryptographic frameworks used?	NO	No custom cryptography. All implementations use standard, well-maintained libraries: argon2-cffi (Argon2id), cryptography (MultiFernet), hashlib (SHA-256), and secrets module for token generation.
PRO-28	What is your key management approach?	PARTIAL	Encryption keys stored as environment variables via Kubernetes Sealed Secrets. MultiFernet versioned key rotation implemented for seamless encryption key transitions without data loss. KMS integration not yet in place.
PRO-29	Do you have a security awareness program?	PARTIAL	People Security Policy documents role-specific security training requirements. Currently sole founder with deep security domain expertise; formal security awareness program framework ready for team scaling.

Reactive Security

5/6 controls satisfied

ID	Question	Answer	Explanation
RS-1	How do you keep aware of potential security vulnerabilities and threats?	YES	Automated scanning with 12 integrated security scanners, EPSS enrichment from FIRST.org API, container image scanning in CI/CD, Falco runtime monitoring, and security alert system with automated threat detection.
RS-2	How do you log and alert on security events?	YES	Structlog JSON logging with 40+ audit event types. Security alert system detects brute force, credential stuffing, API anomalies, and suspicious patterns. Email and Slack notifications for critical findings.
RS-3	Do you have a Security Incident Response Program?	YES	Incident Response Plan (FIMIL-IRP-001) defines four severity levels, incident commander role, response phases (triage, containment, eradication, recovery, communication), and regulatory notification procedures.
RS-4	How is the Incident Response Program tested?	PARTIAL	IRP is documented and technical capabilities are implemented (Incident model, SecurityAlert, auto-blocking, account lockout). DR test completed March 2026 validated recovery procedures. Tabletop exercises and simulated incident drills planned but not yet conducted.
RS-5	Do you have a formal SLA for incident response?	YES	IRP defines response timelines by severity: Critical incidents require immediate response, Major within 1 hour, Minor within 4 hours. SECURITY.md provides 48-hour acknowledgment SLA for vulnerability reports.
RS-6	Do you have formally defined criteria for notifying clients during an incident?	YES	IRP includes customer notification procedures with breach notification aligned to GDPR 72-hour and CCPA timelines. Notification criteria defined by incident severity and data impact.

Software Supply Chain

8/10 controls satisfied

ID	Question	Answer	Explanation
SSC-1	Do you perform static code analysis?	YES	Semgrep and Bandit run in CI/CD on every push and pull request. Ruff linter enforces Python code quality rules. ESLint with strict TypeScript rules (--max-warnings 0).
SSC-2	How do you ensure code is developed securely?	YES	OWASP-aware implementation with Pydantic validation, SQLAlchemy ORM, CSRF protection, and RBAC. Pre-commit hooks enforce code style and secret scanning. CI gates on linting, tests, type checking, and SAST.
SSC-3	Do you perform threat modeling during the design phase?	YES	Formal threat model using STRIDE methodology covering 3 areas with 16 identified threats and mitigations. Risk Assessment (FIMIL-RISK-001) identifies 15 risks with treatment plans. Defense-in-depth architecture reflects threat awareness.
SSC-4	Do you provide developer training in secure coding?	PARTIAL	People Security Policy documents role-specific training requirements including secure coding practices. Currently sole founder with deep security expertise; formal training program ready for team scaling.
SSC-5	What percentage of production code is covered by automated tests?	YES	Frontend enforces 80% coverage thresholds for lines, functions, branches, and statements. Backend has comprehensive test suites including security-specific tests. Exact backend coverage percentage tracked in CI.
SSC-6	Do you have a staging or pre-production system for validating builds?	YES	Clear environment separation: Docker Compose for development, SQLite in-memory for testing, and feature flags for staged rollout. Helm atomic deployments with automatic rollback on failure in production.
SSC-7	Do you maintain a bill of materials for third-party libraries?	YES	Syft SBOM scanner integrated into the platform generates software bill of materials. Poetry (Python) and npm (TypeScript) lockfiles track all dependency versions.
SSC-8	How do you monitor vulnerabilities in third-party dependencies?	YES	Trivy, Grype, and OSV-Scanner provide SCA scanning. EPSS enrichment scores exploit probability. Reachability analysis distinguishes direct from transitive dependency vulnerabilities.
SSC-9	Do you outsource any development?	NO	All development is performed internally by the founder. Vendor Risk Management Policy covers third-party security requirements if outsourcing occurs in the future.
SSC-10	Do you perform security reviews on custom-built software?	YES	Fimil scans its own repositories through the platform. SAST (Semgrep, Bandit) runs in CI. Security-specific tests cover OAuth, token security, webhook signature verification, and rate limiting.

Customer Facing Application Security

11/13 controls satisfied

ID	Question	Answer	Explanation
CFAS-1	How do you authenticate users? What password complexity and SSO options are available?	YES	JWT tokens for API, Redis sessions for web UI. Password policy: 12+ chars, mixed case, digit, special char. Argon2id hashing. TOTP-based MFA available. OAuth2/OIDC SSO with GitHub and generic OIDC providers available on all plans.
CFAS-2	Does the application allow user MFA enforcement by admins?	YES	TOTP-based MFA implemented with recovery codes, two-step login flow, and encrypted secret storage. Users can enable MFA on their accounts.
CFAS-3	Is IP whitelisting available for authentication?	YES	IP blocklist/allowlist functionality with auto-blocking (20+ failed attempts trigger 24-hour block). Admin endpoints for IP management via Security Ops dashboard.
CFAS-4	Are there standardized roles and permissions?	YES	Five standardized roles: Operator (full system), Admin (full tenant access), Security (manage findings/triage), Developer (view findings, limited triage), and Viewer (read-only).
CFAS-5	Are custom granular permissions and roles available?	NO	The platform uses a fixed five-level role hierarchy. Custom granular permissions beyond the predefined roles are not currently supported.
CFAS-6	Are there audit trails and logs for systems with customer data access?	YES	Comprehensive audit logging with 40+ event types tracking actor, tenant, IP, user agent, request ID, and impersonation context. CSV export available for offline analysis.
CFAS-7	Does the application provide admin access to verbose audit logs?	YES	Admin dashboard provides access to audit logs with filtering by user, action type, and date range. CSV export enabled for detailed analysis. API token usage tracked.
CFAS-8	Is a custom data retention policy available for customer data?	YES	Data retention configurable at 30 days for scan reports. Data Governance Policy defines retention schedules for all data categories. Account closure and data deletion procedures documented.
CFAS-9	Does the application provide a change log?	YES	Audit logging captures all CRUD operations and configuration changes with timestamps. Finding status transitions tracked with full history. Version-controlled codebase with Git history.
CFAS-10	Is a sandbox environment available for customer testing?	PARTIAL	Docker Compose development environment available. Enterprise self-hosted model allows customers to run in isolated environments. No dedicated multi-tenant sandbox for SaaS customers yet.
CFAS-11	Is API rate limiting implemented?	YES	Redis-backed sliding window rate limiting: authentication endpoints at 10 requests/minute, general API at 100 requests/minute. Configurable thresholds.
CFAS-12	How do you store API keys?	YES	API tokens are SHA-256 hashed before storage; the plaintext token is shown only once at creation. Tokens are scoped and revocable with audit trail.
CFAS-13	Is IP whitelisting available for API access?	YES	IP blocklist/allowlist functionality available for API access. Admin-managed via Security Ops dashboard with auto-blocking for suspicious activity.

Compliance

4/7 controls satisfied

ID	Question	Answer	Explanation
COMP-1	How do you conduct internal audits?	YES	Comprehensive internal compliance assessment performed against ISO 27001:2022 and SOC 2 Type II criteria with documented findings, gap analysis, and remediation plans. Statement of Applicability tracks control status.
COMP-2	How do you conduct external audits?	NO	No external audit or independent security assessment has been conducted. External penetration testing and certification audits are planned.
COMP-3	Which IT operational, security, or privacy standards do you comply with?	YES	Controls aligned with ISO 27001:2022 and SOC 2 Type II. GDPR and CCPA compliance tracked in Compliance Register. Formal certification not yet obtained.
COMP-4	Do your confidential data access controls align with your classification matrix?	YES	RBAC enforces access by data classification level: Restricted data (encryption keys, tokens) requires Admin/Operator role; Confidential data (PII, scan results) is tenant-isolated with role-based access.
COMP-5	Do you share customer data with any third parties?	PARTIAL	Sub-processors (DigitalOcean, Stripe, Resend, PostHog) may process limited customer data as documented in the sub-processor list. DPAs executed with all vendors. No customer data is sold.
COMP-6	Do you seek the right to use or own customer derived data?	NO	Fimil does not claim ownership of customer data or derived data. Customer data remains the property of the customer as defined in the Terms of Service and DPA.
COMP-7	Is your Privacy Notice externally available?	YES	Privacy Policy published at /privacy. Cookie Policy at /legal/cookies. Data Processing Agreement at /legal/dpa. Acceptable Use Policy at /legal/acceptable-use. All publicly accessible.