Tidela

Security Review Packet

Written answers to the standard vendor security questionnaire, for the IT, security, and procurement teams evaluating Tidela.

> **How to read this document.** Every answer describes what is running in production today. Where a control is not implemented, the answer says so plainly and is marked **Not implemented** — we would rather lose a deal on an honest answer than win one on a claim that falls apart in month two. Tidela has not completed a SOC 2 audit; section L explains our position on that in full.

Product: Tidela — AI Selling platform (CRM replacement)
Vendor: Surf and Sand Consulting, LLC · Florida, United States
Service URL: app.tidela.ai · **Version of this document:** 1.0, August 10, 2026
Security contact: security@tidela.ai

A Company & service overview

A1 **What does the service do?**

Tidela is a multi-tenant, cloud-hosted B2B sales platform: accounts, contacts, opportunities, activities, projects, and quotes, with AI assistance layered throughout. It replaces a traditional CRM.

A2 **Who is the legal entity?**

Surf and Sand Consulting, LLC, a Florida limited liability company. Tidela is its product.

A3 **How is the service delivered?**

Software as a service, over HTTPS, in a browser. There is also an installable progressive web app, a REST API, and an MCP server for AI-tool integration.

A4 **Do you have a named security contact?**

Yes — security@tidela.ai, monitored by the engineering lead. Acknowledgement within two business days.

A5 **How large is the team with production access?**

Production access is limited to the founder/engineering lead. **We state this plainly because it cuts both ways:** the access surface is minimal and there is no offshore support tier touching your data, but there is also no separation-of-duties between developer and operator. Section K describes the compensating controls.

A6 **Do you use subcontractors or offshore development?**

No. Development is in-house in the United States. Third-party services that process data are listed in section J.

B Architecture & hosting

B1 **Where is the application hosted?**

DigitalOcean, United States region. DigitalOcean maintains its own platform-level compliance certifications; those cover the infrastructure layer only, not Tidela's application controls.

B2 **What is the technology stack?**

Python/Django application served by gunicorn in a Docker container behind Caddy (automatic TLS), with PostgreSQL for data and Redis for cache and sessions.

B3 **Is customer data stored outside the United States?**

No. All customer data is stored in the United States.

B4 **Do you support regional data residency (EU, Canada)?**

Not implemented. The data model is region-aware — identity and audit records carry a region — so residency is a deployment exercise rather than a rewrite, but infrastructure is single-region today. Tell us before signing if this is a requirement.

B5 **Is the environment dedicated or shared?**

Shared, multi-tenant, with logical isolation enforced in the application and data layer (section C). Dedicated single-tenant deployment is not offered.

B6 **Is there a separate non-production environment?**

Development and automated testing run locally against a separate database, never against production data. A dedicated hosted staging environment is on the near-term roadmap.

B7 **Is administrative infrastructure access protected by MFA?**

Yes — hosting provider console access requires multi-factor authentication; server access is by SSH key, not password.

C Multi-tenancy & data isolation

C1 **How is one customer's data separated from another's?**

Every record belongs to exactly one workspace, and every query is scoped to the requesting user's active workspace in application code. Workspace membership is validated on each request against the user's memberships, not taken from client-supplied input.

C2 **Can a user belong to more than one workspace?**

Yes, by explicit invitation. There is no automatic tenant join — a matching email domain never grants access to an existing workspace.

C3 **Has isolation been tested?**

Yes. A dedicated security review in July 2026 traced tenant isolation end to end across the application, the API, and the AI layer, and found no cross-tenant access paths. It is re-tested as part of the automated test suite.

C4 **Are object identifiers guessable, and does that matter?**

Identifiers are sequential, and access is authorised on every request regardless — knowing another workspace's record ID does not grant access to it.

C5 **Are backups isolated per customer?**

No — backups are taken at the database level and contain all workspaces. They are encrypted and access is restricted to the same single administrator.

D Identity & access management

D1 **Do you support single sign-on?**

Yes — SAML 2.0 and OpenID Connect, with Okta, Microsoft Entra ID, Google Workspace, and any standards-compliant provider. SAML supports both SP-initiated and IdP-initiated sign-in.

D2 **Is SSO an extra-cost feature?**

No. There is no per-seat security fee and no price premium for connecting your identity provider. Enterprise Security is a per-workspace setting Tidela enables on request at no additional charge; it is off for a new self-serve workspace, and enabling it activates role enforcement, the MFA policy, SCIM, and the audit trail. Ask for it at the start of an evaluation so you are testing the configuration you would actually run.

D3 **Can SSO be enforced for all users?**

Yes. Workspace owners retain password sign-in as a deliberate break-glass path so a misconfigured IdP connection cannot lock an organisation out of its own data.

D4 **Do you validate SAML assertion signatures?**

Yes, and assertion signing is required by default. Tidela does not sign its own authentication requests and does not require encrypted assertions — TLS protects the transport — so there is no service-provider keypair for your team to manage.

D5 **Do you support SAML Single Logout?**

Not implemented. Deprovisioning through SCIM revokes sessions immediately, which covers the material risk.

D6 **Do you support SCIM provisioning?**

Yes — SCIM 2.0, with create, activate, and deactivate. A bearer token is generated per workspace and stored hashed. Note: attribute updates pushed from the IdP (name, email changes) are accepted but not currently written back to the user record — only the active/inactive state is applied.

D7 **What happens when we deprovision a user in our IdP?**

Their Tidela account is deactivated and their active sessions and API/OAuth tokens are revoked at the same time. Access ends immediately, not at next login.

D8 **Can IdP groups drive permissions?**

Yes. Map an IdP group to a Tidela role and role assignment follows your directory, over both OIDC and SAML. Users without a mapped group receive a configurable default role.

D9 **Do you support multi-factor authentication?**

Yes — authenticator-app TOTP with one-time recovery codes. Any user can enable it; admins can require it workspace-wide, and unenrolled members are gated until they enrol. SSO users' factors are handled by your IdP.

D10 **Do you support SMS-based MFA?**

No, deliberately. SMS is vulnerable to SIM-swap; TOTP is not.

D11 **Do you support passkeys or hardware keys?**

Not implemented in Tidela directly. Available today by fronting Tidela with an IdP that supports them and enforcing SSO.

D12 **Describe your authorisation model.**

Role-based. Each role carries per-module permissions (view, create, edit, delete, export) across every module, plus a record scope of all, team, or own. A role with no grant for a module gets nothing. Note: administrators re-permission the seeded roles (owner, admin, manager, rep, read-only) rather than creating additional named roles; role creation in the UI is not implemented.

D13 **Are permissions enforced consistently outside the UI?**

Yes — the web application, the REST API, the MCP tool layer, and the AI assistants all call one shared server-side permission check, not parallel implementations that drift. This includes the workspace-level assistant: the grounding data it reasons over is narrowed to the asking user's record scope, so it cannot summarise records that user could not open in the app.

D14 **What password policy applies to local accounts?**

Django's standard validators: minimum length, similarity to user attributes, common-password blocklist, and numeric-only rejection. Alignment to NIST SP 800-63B, including an admin-configurable minimum length and a breach-corpus check, is on the near-term roadmap.

D15 **Is there brute-force protection on sign-in?**

Yes — failed attempts are throttled with a temporary lockout, and the client IP is derived from a trusted-proxy count so it cannot be spoofed by a forwarded header.

D16 **How are sign-in links handled?**

Magic links are single-use, short-lived, and bound to a token identifier that is burned on redemption. Credentials never appear in a URL.

D17 **Can administrators see user passwords?**

No. Passwords are stored only as salted hashes and are never recoverable.

D18 **Can users view and terminate their own sessions?**

Not implemented as a self-service screen today; it is the next item on our roadmap. Sessions can be revoked now by deactivating the user, via SCIM or in the admin.

E Encryption & key management

E1 **Is data encrypted in transit?**

Yes — HTTPS/TLS on every connection, with certificates provisioned and renewed automatically. HTTP is redirected, and HSTS is set.

E2 **Is data encrypted at rest?**

Yes, at the storage layer provided by our hosting platform.

E3 **Are credentials and secrets encrypted separately?**

Yes. Integration tokens, MFA secrets, SSO client secrets, and platform API keys are encrypted at the application layer with Fernet (AES-128-CBC with HMAC authentication) before they reach the database, and are never rendered back into a page.

E4 **Where do encryption keys live?**

In the server environment configuration, never in source control. Key rotation is manual and documented.

E5 **Do you support customer-managed encryption keys?**

Not implemented.

E6 **Are API tokens stored in plaintext?**

No — API keys and SCIM tokens are stored as hashes and shown to the administrator exactly once at creation.

F Application security

F1 **How do you prevent SQL injection?**

All data access goes through Django's ORM with parameterised queries. There is no string-built SQL in the application.

F2 **How do you prevent cross-site scripting?**

Template auto-escaping throughout, a dedicated safe serialiser for embedding data in scripts, and a strict allowlist sanitiser for rendering synced email HTML. A stored-XSS finding identified in a June 2026 review was fixed and verified.

F3 **Is CSRF protection in place?**

Yes, on all state-changing requests.

F4 **What security headers do you set?**

Content-Security-Policy, Permissions-Policy, HSTS, X-Content-Type-Options, X-Frame-Options, and Referrer-Policy in production.

F5 **How do you handle server-side request forgery?**

All outbound fetches of user-supplied URLs go through a single guarded fetcher that resolves and pins the destination address, blocks loopback, link-local, cloud-metadata, and private ranges, and refuses redirects into internal hosts. An SSRF issue found in the July 2026 review was fixed and confirmed in production.

F6 **Are file uploads restricted?**

Uploads are stored as attachments outside the application code path and served with a non-executing content disposition.

F7 **Do you conduct security reviews?**

Yes. Structured internal reviews covering isolation, authentication and authorisation, the AI attack surface, and classic web and configuration risk were completed in June and July 2026; findings were remediated and the fixes verified in production. Reports are available under NDA.

F8 **Do you commission third-party penetration tests?**

Not yet. Reviews to date have been internal and structured, not an independent third-party test. This is the honest gap in our application-security story and we will say so rather than imply otherwise.

F9 **Do you have a vulnerability disclosure policy?**

Yes — published at tidela.ai/enterprise-security and machine-readable at /.well-known/security.txt. Good-faith research within the stated bounds is welcome and will not be pursued.

F10 **How do you manage dependency vulnerabilities?**

Dependencies are pinned with floors that exclude known-vulnerable versions and are updated on review. Automated dependency scanning in the build pipeline is being added now.

F11 **Is there a test suite?**

Yes — approximately 1,000 automated tests covering permissions, isolation, and business logic, run before each deployment.

F12 **How are changes deployed?**

Container image rebuild and restart, with database migrations applied automatically at boot and a restore point captured before each deploy so a release can be rolled back. Migration to a CI-driven pipeline with a staging gate is in progress.

G AI & data handling

The section most vendor questionnaires have not caught up to yet, and the one we consider most important to answer precisely.

G1 **Which AI provider do you use?**

Anthropic's Claude models over their commercial API, by default. A workspace may supply its own key for another provider (OpenAI, Google, or an OpenAI-compatible endpoint); if it does, prompts go to that provider under your agreement with them. This is off unless you enable it.

G2 **Is our data used to train models?**

No. Tidela does not train models on customer data and does not permit its AI provider to use customer data for model training under the commercial terms it operates under.

G3 **What data is sent to the model, and when?**

The specific records relevant to the action a user just took — the account, contact, deal, or message in front of them. AI runs on demand when a user clicks, plus the scheduled jobs you enable yourself (weekly digest, account briefs). There is no continuous background sweep of your database. Research and enrichment features additionally consult public web sources and the prospect's own website, and indicate this on screen.

G4 **Can the AI reach data the user cannot?**

No. AI features execute inside the same workspace scoping and role record scope as the invoking user — including the workspace-level assistant, whose grounding snapshot is filtered by that user's permissions before it reaches the model.

G5 **Can the AI take actions on its own?**

Only additive, reversible ones, and only from a fixed server-side allowlist — creating a note, a task, or a draft. An assistant never sends email; it drafts for a person to send. Where a human has built an email sequence or enabled the weekly digest, Tidela sends on the schedule that person approved. Deletions and configuration changes are never AI-initiated.

G6 **How do you handle prompt injection?**

Authorisation decisions are made in application code, never delegated to model output. Content the model reads — an email body, a synced message, a fetched web page — is treated as untrusted data, so an instruction hidden inside it cannot widen the assistant's permissions or trigger an unallowlisted action.

G7 **Does the AI invent facts about our pipeline?**

Numbers that must be exact — forecasts, weighted pipeline, renewal dates — are computed in application code, not generated by the model. Where the model offers an external or firmographic claim it is surfaced as unverified rather than asserted.

G8 **Is AI usage metered or capped?**

Yes — daily cost and call-rate caps are enforced ahead of model calls, with request timeouts, so a loop or abusive script is bounded rather than open-ended. Caps are enforced on the interactive AI request path; workspaces using their own AI key are billed by their own provider and are subject to that provider's limits rather than ours.

G9 **What does your internal usage analytics store?**

Event counts and metadata for capacity and product analytics. It does not store prompt text or record contents.

G10 **Can AI features be turned off?**

Yes — AI can be disabled at the workspace level.

G11 **Does Tidela track opens and clicks on email we send?**

Only if you switch it on. Open/click tracking is **off by default** for new workspaces: no pixel, no link wrapping, until an admin enables it under Admin → Organization. We changed this default deliberately — tracking the people your team emails should be an explicit decision, not an inherited one. Workspaces created before August 2026 may already have it enabled; it is visible and reversible in that same screen.

H Logging, monitoring & audit

H1 **What is captured in the audit log?**

Authentication events; users created, invited, deactivated and reactivated; role assignments; team creation, deletion and membership changes; IdP group→role mappings; API keys created, rotated and revoked; AI-assistant (OAuth) connections revoked; MFA and SSO configuration changes; data exports; administrative and platform-operator actions; partner-portal sign-ins and submissions; and the Enterprise Security entitlement itself being switched on or off. Each entry records the actor, action, target, before/after values where applicable, IP address, and UTC timestamp.

H2 **Can we search and export the audit log?**

Yes — filter by category, actor, and time in the application, and export to CSV or JSON.

H3 **Can we stream audit events into our SIEM?**

Export today is administrator-initiated CSV/JSON. A read-only audit-events API endpoint for scheduled SIEM collection is in active development; push-based streaming is not planned near-term. **We regard getting these events into a system we do not control as the strongest integrity guarantee available** and we will help you wire it up.

H4 **Is the audit log tamper-resistant?**

Entries are chained with an HMAC, so alteration or silent deletion by an actor with database access but not the application key is detectable, and the application reports chain verification status. **Scoped honestly:** the key is reachable by the application, so this does not defend against a full compromise of the application itself. Database-enforced append-only permissions are a near-term hardening step, and SIEM export (H3) remains the real answer.

H5 **Can we set audit retention?**

Yes, per workspace; the default is 365 days.

H6 **Are field-level record changes tracked?**

Yes, on each record's own history, separately from the security audit log.

H7 **Is application error monitoring in place?**

Yes, with structured application logging and an error-monitoring service (section J).

H8 **Is there uptime monitoring and a public status page?**

A health endpoint is checked on deploy. External uptime monitoring with a public status page is being stood up now.

I Data lifecycle

I1 **Who owns the data?**

You do. Tidela claims no ownership of customer data and does not sell it or share it with advertisers or data brokers.

I2 **Can we export our data?**

Yes — CSV export for accounts, contacts, deals, and projects, plus a full REST API with keyset pagination (the API is a separately enabled capability). CSV export covers standard fields; custom fields, activities, and attachments come out through the API. A single-click whole-workspace export covering everything is in development.

I3 **Is export permission-controlled?**

Yes. Export is a distinct permission per module per role, and every CSV export path — entity exports, saved reports, the account ad-audience export, partner payout reports and import error reports — is permission-gated and writes a `data.export` audit event naming what left and who took it.

I4 **What happens on termination?**

Data is returned or deleted at your election. Absent a contrary instruction, workspace data is deleted from live systems within 30 days and from backups within the ordinary rotation thereafter.

I5 **Does deleting a workspace remove everything?**

Yes, including the company and user records mirrored into our own internal Tidela workspace for trial and support context — specifically implemented and verified rather than assumed. Data held by operational subprocessors (error monitoring, email delivery logs) ages out on their own retention schedules.

I6 **Can we set retention rules for records?**

Audit retention is configurable now. Configurable retention for synced email and activity records is on the roadmap; individual records can be deleted at any time.

I7 **Are deletions recoverable?**

Records are archived first and restorable by an administrator, then hard-deleted on request — so an accidental deletion is recoverable and an intentional one is final.

J Subprocessors

J1 **Who are your subprocessors?**

Core: DigitalOcean (hosting, database, storage — US); Anthropic (AI models — US); Postmark (system and outbound mail — US). **Operational:** Sentry (error monitoring); Cloudflare Turnstile (bot protection on sign-up and sign-in-link requests); OpenStreetMap (address geocoding via Nominatim, plus map tiles, which see the viewer's IP and the map area); Clearbit, Google's favicon service and DuckDuckGo's icon service (company logo lookup by website domain); Gravatar (contact avatar lookup by hashed email address); Google Fonts and unpkg (content delivery for the typeface and the map/org-chart libraries); browser push services (Apple, Google, Mozilla) where a user enables notifications. **Optional, only when a customer connects them:** Google (Gmail/Calendar sync); Microsoft (Microsoft 365 sync); ElevenLabs (narration of Tidela's own training content — receives no customer data); a previous CRM (HubSpot, Pipedrive, Attio, Nimble, Copper, Capsule, monday.com, Pipeliner) when you run a migration import; and an alternative AI provider if you supply your own key. The current list is maintained at tidela.ai/legal/subprocessors.

J2 **Will we be notified of new subprocessors?**

Yes, on request — email security@tidela.ai to be added to the notification list. Notice is given before a new subprocessor that handles customer records begins processing.

J3 **Do subprocessors have contractual data protection obligations?**

Yes, no less protective than those in our own DPA, and Tidela remains responsible for their performance.

K Business continuity & incident response

K1 **Are backups taken?**

Yes, automated at the database platform level, encrypted, and retained on a rolling schedule.

K2 **What are your RPO and RTO?**

We are not publishing numbers we have not measured. A documented restore drill is scheduled, and tested recovery-point and recovery-time figures will be published here and on our security page once we have them. Ask for the current status — we will give you the real answer.

K3 **Do you offer a contractual uptime SLA?**

Not today. We would rather add one alongside the monitoring and redundancy that make it meaningful than print a percentage we cannot evidence. Available for discussion on an enterprise agreement.

K4 **Is the service highly available?**

The application runs multiple worker processes on a single host with a managed database. There is no multi-node failover today; recovery from host failure is a restore, not an automatic cutover. This is the clearest trade-off of our current stage and we state it directly.

K5 **Do you have an incident response process?**

Yes — a documented runbook covering detection, containment, assessment, customer notification, remediation, and post-incident review. Given team size, escalation is direct rather than tiered, which in practice means the person who can fix it is the person who is paged.

K6 **Will you notify us of a breach, and how fast?**

Yes — without undue delay and within 72 hours of becoming aware, with what is known at the time and updates as the investigation proceeds. This is a contractual commitment in our DPA, not a courtesy.

K7 **Do you carry cyber liability insurance?**

Ask us for the current position before contracting; we will answer in writing rather than in a marketing document.

K8 **What happens to our data if the company ceases operations?**

A fair question to ask a small vendor. Your data is exportable at any time in open formats, with no proprietary lock-in, and we will agree contractual notice and an export window on an enterprise agreement.

L Compliance & certifications

L1 **Are you SOC 2 certified?**

No. Tidela has not undergone a SOC 2 audit, and we do not describe ourselves as "SOC 2 aligned" or "SOC 2 ready" — phrases that mean nothing without an auditor's opinion behind them. Our current position: we have adopted the underlying controls and documented policies, and we will commission an audit when a customer commitment justifies it. If SOC 2 is a hard gate for your organisation, tell us during evaluation — it moves our roadmap.

L2 **Are you ISO 27001 certified?**

No.

L3 **Do you have written security policies?**

Yes — information security, access control, secure development and change management, incident response, business continuity, vendor management, and data classification and retention. Available for review under NDA.

L4 **Will you sign a DPA?**

Yes. Our standard DPA is published at tidela.ai/legal/dpa and we sign it as written for most customers. It covers processor obligations, security measures, subprocessors, breach notification within 72 hours, assistance with data subject rights, and deletion on termination.

L5 **Can you support GDPR and UK data protection requirements?**

Yes as a processor, including Standard Contractual Clauses or the UK International Data Transfer Addendum for transfers to the United States. Data subject requests are largely self-service through Tidela's own search, edit, export, and delete functions.

L6 **Can you support CCPA/CPRA requirements?**

Yes as a service provider. Tidela does not sell or share personal information as those terms are defined.

L7 **Is Tidela suitable for regulated data — PHI, cardholder data, government identifiers?**

No, and we will tell you that up front. Tidela is a B2B sales platform for business contact and pipeline data. We do not sign BAAs and Tidela is not in scope for HIPAA or PCI DSS. Do not put that data in it.

L8 **Will you complete our security questionnaire?**

Yes. This document answers most of them; send anything it misses to security@tidela.ai and you will get written answers from the person who wrote the code, usually within a few business days.

Summary of known gaps

Collected here so nothing in this document has to be hunted for: no SOC 2 or ISO 27001 certification (L1, L2); no third-party penetration test to date (F8); no published uptime SLA and no multi-node failover (K3, K4); RPO/RTO not yet measured (K2); no self-service session management screen (D18); role creation in the UI is not implemented (D12); SIEM export by scheduled API pull in development, no push streaming (H3); single region, no EU or Canadian residency yet (B4); no customer-managed encryption keys (E5); no passkeys native to Tidela (D11); no SAML Single Logout (D5); not suitable for PHI, cardholder data, or government identifiers (L7). Several of these are in active development; ask for current status.

Tidela — Security Review Packet, version 1.0, August 10, 2026. Answers describe the production service as of this date and are provided in good faith for vendor evaluation; they do not constitute a warranty or amend any agreement between the parties. For the current version, corrections, or anything this document does not cover: security@tidela.ai · tidela.ai/enterprise-security