SECURITY AND TRUST WHITEPAPER

Security, privacy and trust at Linkence.

Linkence AI answers questions and automates work over the tools your team already uses: mail, documents, project trackers, chat, CRM and calendars. This whitepaper describes, in detail, how customer data is isolated, encrypted, governed, retained and deleted; how our AI processing works; who our subprocessors are; and how we respond to incidents. It is written for security, IT and procurement teams evaluating Linkence. Contact: security@linkence.ai.

OUR COMMITMENT

Your data is never used to train AI models. Every organization runs in its own isolated tenant, stored credentials are encrypted with AES-256-GCM, outbound agent actions require explicit user approval, and all data is deleted on request with a verifiable deletion ledger.

1. Answers at a glance

QUESTION	ANSWER
Is our data used to train AI models?	No, never. Default OpenAI calls disable provider-side storage; you can also bring your own model keys.
Where is our data hosted?	United States (Railway application and database hosting, Cloudflare edge and file storage).
Is our data isolated from other customers?	Yes. Per-organization databases, per-organization vector collections, and organization-scoped routing.
How is data encrypted?	TLS 1.2 or higher in transit; encrypted at rest; credentials get an extra AES-256-GCM layer.
How do users sign in?	Google, Microsoft, Azure Active Directory and generic OAuth 2.0 providers, or email-based login.
Can we delete our data?	Yes. Hard deletes at document, connector, user or organization level, recorded in a deletion ledger.
How fast do you notify us of a breach?	Within 48 hours of becoming aware, per our Data Processing Agreement.
Have you been penetration tested?	Yes. Independent web-application VAPT completed April 2026 with no critical findings; report available under NDA.
Is a DPA available?	Yes, ready for signature, with a published subprocessor list and 30 days' change notice.
Are you certified?	Not yet. A complete ISO/IEC 27001:2022 and SOC 2 aligned security program is documented and certification readiness is in progress.

2. Architecture and hosting

Linkence is a multi-tenant SaaS platform. Customer traffic terminates at the Cloudflare edge over HTTPS and is routed to the correct tenant's backend based on the organization identity in the request. The application layer consists of an API service, background workers, and a dedicated AI service that performs retrieval and agent orchestration. Services are hosted on Railway in the United States.

STORE	PURPOSE	PROTECTION
MongoDB	Users, organizations, connector configuration, chat history, agent configuration	Encrypted at rest; per-tenant databases; daily automated backups
Qdrant	Vector search over indexed content	One isolated collection per organization
Redis	Queues, sessions, caching, approval round-trips	Managed hosting; transient workflow state only
Cloudflare R2	Uploaded files and generated artifacts	Object encryption at rest and provider durability

How your data flows through the platform

Connect
Your admin authorizes each connector via OAuth with scoped permissions. Tokens are encrypted before storage.

→

Sync & index
Content is synced over HTTPS and indexed inside your organization's own isolated collection.

→

Retrieve
A user question retrieves only the relevant excerpts from your data, scoped to your organization.

→

Infer
Excerpts are sent to the model over HTTPS. No training, and no provider-side storage on the default path.

→

Answer
The user receives a grounded answer with citations back to the source documents.

What leaves your systems, and what does not

- **Indexed content:** documents and messages from connectors you enable are synced into your isolated tenant so they can be searched. You choose which connectors and scopes to enable.
- **Live-fetch mode:** for sensitive classes of data, content can be fetched live from the source at question time instead of being stored, which removes at-rest copies entirely.
- **Model calls:** only the excerpts retrieved for a specific question are sent to the model, never your bulk corpus.
- **Internal traffic:** service-to-service calls are authenticated with an internal API key plus short-lived signed service tokens.

3. Tenant isolation

- **Isolated data stores:** each organization has its own database namespace and its own dedicated vector collection. Cross-tenant queries are not possible by construction; every search carries organization-scoped access filters.
- **Organization-scoped routing:** the edge layer resolves the organization identity on every request and routes it to that tenant's backend. Requests without a valid organization context never reach tenant data.
- **Scoped request context:** organization, user and request identifiers are carried through every internal call and every log line, so all processing is attributable to exactly one tenant.
- **Permission-aware retrieval:** retrieval respects team and role scoping inside your organization, so users only search content they are entitled to see.

4. Encryption and key management

In transit and at rest

- All customer-facing traffic is served over HTTPS with TLS 1.2 or higher, terminating at the Cloudflare edge. All provider and webhook traffic is HTTPS only.
- Primary data stores (MongoDB, Redis, Qdrant, Cloudflare R2) are encrypted at rest by the managed hosting providers.
- High-sensitivity values get an additional application-layer AES-256-GCM encryption before storage: connector OAuth tokens, database connection credentials, customer model API keys, SMTP credentials, and identity provider client secrets and certificates.

Approved cryptography

USE	PRIMITIVE
Stored credentials	AES-256-GCM authenticated encryption with a unique random IV per operation
Passwords	bcrypt, cost factor 12
Sessions and signed URLs	Signed JWTs with separate secrets per token class; 30-minute access tokens
Reset and invite tokens	SHA-256 hashes of 32-byte random tokens; never stored in plaintext
Inbound webhooks	HMAC-SHA256 signature verification with constant-time comparison

Weak or legacy algorithms (MD5, SHA-1, DES, RC4, ECB mode) and home-grown cryptography are prohibited by policy.

Key handling and rotation

- Production keys exist only in the hosting provider's environment configuration; they are never committed to source control, and repository history is verified clean of secrets.
- Secrets are never written to logs; structured loggers redact secret-bearing fields before emission.
- Keys are rotated at least every 12 months, within 5 business days after personnel departure, and immediately on suspected compromise. Every rotation is recorded as audit evidence.

5. Identity, authentication and access control

Your users

- **Organization sign-in:** organization-level sign-in via Google, Microsoft, Azure Active Directory and generic OAuth 2.0 providers. Identity provider client secrets are encrypted at rest.
- **Sessions:** short-lived signed access tokens (30 minutes) with separate refresh tokens (6 hours); passwords are hashed with bcrypt.
- **Role-based access control:** admin and member roles per organization. Only admins can manage members, organization settings, model settings, analytics and deletion operations. Users cannot promote themselves, and the last remaining admin cannot be demoted.
- **Audit trail:** role changes are recorded with who changed what and when, and written to the customer-visible audit trail. Cached authorization data is re-validated against the database on a bounded interval.

Our personnel

- **Least privilege:** every internal account receives the minimum access required for its documented function.
- **Quarterly access reviews** cover infrastructure consoles, repository access, service credentials and any standing production access; results are retained for at least 3 years.
- **MFA on infrastructure consoles** (hosting, edge, source control) is mandatory and verified at each review.
- **Deprovisioning:** all internal access is revoked within one business day of departure, same day for involuntary terminations, with associated key rotation.
- **Production data access** is exceptional, justified, time-bound and logged; routine operations use application functionality or reviewed scripts.

6. Connector and integration security

- **Scoped OAuth consent:** every connector is authorized by your admin through the provider's OAuth flow with defined scopes. You choose which connectors and which scopes to enable, and the requested permissions are visible at consent time.
- **Token protection:** access and refresh tokens are encrypted with AES-256-GCM before storage and are never logged or returned to the browser.
- **Revocation:** disconnecting a connector cancels in-flight jobs, deletes the synced data, revokes the OAuth grant at the provider, and removes webhook routes.
- **Webhook authentication:** inbound provider webhooks are verified with HMAC-SHA256 signatures using constant-time comparison before any processing.
- **Outbound actions:** agent actions that leave the platform, such as sending an email, require explicit user approval before execution and fail closed if their inputs cannot be verified.

7. AI and model governance

> **THE CORE GUARANTEE**
>
> **Customer content is never used to train or fine-tune AI models**, neither by Linkence nor by our model providers on the paths we control.

Model providers

PROVIDER	ROLE	DATA PROTECTION
OpenAI	Default inference and embeddings	API calls made with provider-side storage disabled, enforced in platform code
Anthropic	Optional, customer-selectable	Zero-data-retention configured at the provider account level
Google Gemini	Optional, customer-selectable	Zero-data-retention configured at the provider account level
Bring your own model	Your keys, your endpoint	Inference runs under your own provider contract; your endpoints are not our subprocessors

What the model sees

- Model calls carry only the excerpts retrieved for the specific question being asked, assembled at query time, never your bulk corpus.
- Retrieval is scoped to your organization and respects team and role scoping inside it.
- Answers are grounded in retrieved sources and cite them, so users can verify every claim against the original document.

AI data lifecycle

- **Prompts and chat history:** deleted automatically after 15 days; users can delete on demand at any time.
- **Embeddings and indexed content:** retained in your isolated tenant until you delete the document, disconnect the connector, remove the user, or delete the organization.
- **Agent working state:** held in process memory for the lifetime of a request and not durably persisted.

Human control over agent actions

- Actions with outside effect, such as sending email or posting to connected tools, require explicit user approval before execution.
- If a requested attachment or input cannot be verified, the action fails closed and tells the user exactly what was not executed, rather than guessing.
- A data protection impact assessment covering AI processing has been completed and is available on request.

8. Data retention and deletion

Retention is enforced automatically by store-level expiry indexes and by on-demand deletion. Expiry dates are stamped when data is written, so retention does not depend on batch jobs.

DATA CLASS	RETENTION
Chat sessions and messages	15 days, deleted automatically; users can delete on demand at any time
Automation events (auto-reply and similar)	30 days, deleted automatically
Skill run records	90 days, deleted automatically
Production error records	90 days after last occurrence; contains error metadata, not document content
Synced documents, chunks and embeddings	Until you delete the document, disconnect the connector, remove the user, or delete the organization
OAuth tokens	Hard deleted and revoked at the provider on disconnect or account deletion
Deletion ledger	Retained as audit evidence of deletions performed

Deletion on request

- **Single document:** removes stored rows and vector points, and writes a tombstone so a later connector sync cannot resurrect the deleted item.
- **Connector disconnect:** cancels in-flight jobs, deletes all synced data, revokes the provider OAuth grant, and removes webhook routes.
- **User removal:** purges the user's documents and personal connectors, with provider-side token revocation.
- **Organization deletion:** revokes all organization connector grants, purges all organization data, and drops the organization's vector collections.
- **Verification:** every deletion is recorded in a control-plane ledger with status, deleted counts and residue verification, available as audit evidence.

No secrets or content in logs

- Structured loggers on every service redact secret-bearing and content-bearing fields by denylist before anything is emitted: passwords, tokens, credentials, message content, prompts and email addresses.
- Agent audit logging strips document content from centralized logs; full detail remains visible only to the requesting user's own session.

9. Secure development and vulnerability management

Continuous security scanning

LAYER	TOOLING AND CADENCE
Static code analysis	CodeQL on every pull request and weekly
Dependencies	npm audit and pip-audit on every pull request and weekly; Dependabot patches
Containers	Trivy image and filesystem scans on every pull request and weekly
Secrets	Gitleaks on every pull request and weekly; committed secrets are treated as incidents

- Findings are triaged within one business day, assigned CVSS-based severity adjusted for real exploitability, and remediated under severity-based deadlines; if a full fix cannot ship in time, an effective mitigation is deployed within the same deadline.
- Findings suggesting active exploitation or customer data exposure escalate immediately to incident response.
- **Independent penetration testing:** a third-party web-application VAPT was completed in April 2026 with no critical vulnerabilities identified; lower-severity findings were remediated or formally risk-accepted. Testing repeats at least annually and after major architectural changes. The report is available to customers under NDA.

10. Availability, backup and disaster recovery

- **Recovery objectives:** service restoration within 24 hours and maximum data loss of 24 hours, reviewed against periodic restore tests. We never promise better than tested capability.
- **Backups:** all databases are backed up daily, encrypted, and stored independently of the primary hosting failure domain. Uploaded files rely on provider object durability.
- **Re-derivation:** indexed content can additionally be rebuilt by re-syncing your connectors, which materially reduces permanent-loss risk.
- **Monitoring:** continuous production error detection with automated alerts; unhealthy deployments are detected within minutes and rolled back.
- **Exercises:** business continuity is exercised at least annually, including a real restore into a non-production environment, with documented results.

11. Incident response and breach notification

- **Detection:** automated production error monitoring, health and deploy alerts, CI security gates, customer reports and provider notifications all feed a single documented incident process with severity levels.
- **Assessment within 24 hours:** every suspected personal data breach is formally assessed within 24 hours of detection, including scope, affected data subjects and risk level. The assessment is recorded even when the conclusion is no breach.
- **Customer notification within 48 hours:** where Linkence processes data on your behalf, you are notified without undue delay and in any case within 48 hours of awareness, with the facts, likely consequences and measures taken, followed by updates until resolution and a written post-incident summary.
- **Regulatory support:** notifications support your own obligations to supervisory authorities; where Linkence acts as controller, authorities are notified within 72 hours.
- **Breach register:** every breach is documented with facts, effects and remedial action, regardless of notification outcome.

12. Compliance posture

We state our posture honestly: **Linkence is not yet SOC 2 or ISO/IEC 27001 certified.** What we have built is a complete, documented information security management system aligned to ISO/IEC 27001:2022 and the SOC 2 Trust Services Criteria, and certification readiness is in progress.

- **Documented security program:** 24 approved policy and register documents covering information security governance, risk management, access control, cryptography, secure development, vulnerability management, logging and monitoring, incident response, business continuity, retention, supplier management, HR security and acceptable use. Available for review under NDA.
- **GDPR readiness:** a Data Processing Agreement ready for signature, records of processing activities, a completed data protection impact assessment for AI processing, and documented data subject rights support.
- **Evidence-based:** each control maps to concrete evidence in code and operational records through a maintained evidence matrix, so auditors and customers can verify claims rather than take them on trust.
- **Independent validation:** third-party penetration testing completed with no critical findings, repeated at least annually.

13. Subprocessors

All subprocessors operate under data processing agreements with appropriate transfer safeguards. We give at least 30 days' written notice before engaging a new subprocessor that processes customer personal data, and you may object per the DPA. Feature-specific providers only process data when you enable the corresponding feature.

SUBPROCESSOR	PURPOSE	LOCATION
Railway	Application and database hosting	United States
Cloudflare	Edge routing, CDN, file and artifact storage	Global edge, US jurisdiction
Upstash	Managed Redis for sessions, queues and caching	United States
OpenAI	Default LLM inference and embeddings, storage disabled	United States
Anthropic	Optional LLM inference	United States
Google (Gemini API)	Optional LLM inference and embeddings	United States / global
Unipile	LinkedIn integration (if enabled)	European Union
Twilio	Voice calling (if enabled)	United States
Meta	WhatsApp Business messaging (if enabled)	United States / global
Apollo.io	Lead enrichment (if enabled)	United States
Tavily	Web search for research features (if enabled)	United States
Firecrawl	Scraping of customer-specified public URLs (if enabled)	United States
GitHub	Source code and CI; no customer content	United States

The complete, versioned subprocessor register with data categories and safeguards per provider is published to customers and available at any time.

14. Working with our security team

- **Security questionnaires:** we complete customer questionnaires from our documented control set, typically within one week.
- **Available on request:** Data Processing Agreement, full subprocessor register, penetration test report (under NDA), and the complete policy pack (under NDA).
- **Vulnerability reports:** researchers and customers can report suspected vulnerabilities directly to the address below; reports are triaged within one business day.

SECURITY CONTACT

For security reviews, questionnaires, our DPA and subprocessor list, or to report a vulnerability, write to security@linkence.ai.