Skyseat · Security sheet

Skyseat lets a company run AI employees, each with its own cloud computer, its own permissions and a named human in charge. Skyseat is in early access. This sheet says what exists today and, in the red box, what does not.

WHERE DATA LIVES

- Platform, database and vault: Hetzner, **Helsinki, Finland**. AI employees' computers: Hetzner, Germany. Hetzner holds ISO 27001; we do not operate data centers.
- A **U.S. region** is built behind a switch and will be created when the first U.S. customer needs it.
- Each company's data is separate. An automated test tries every dashboard route against another company's data (all return not-found); new routes join it automatically.
- The AI model (Claude, Anthropic) and the voice/phone providers are in the **U.S.** whatever the region; subprocessors are public with 30 days' notice of changes.

ENCRYPTION AND KEYS

- HTTPS/TLS for all traffic to the platform.
- Customer credentials for connected services: AES-256-GCM in a vault; the key is kept outside the database. An AI employee's computer never sees them: the platform makes the call, checks the permission and logs it.
- Passwords: scrypt, common passwords rejected. Nobody sets a password for someone else (single-use invite links).
- No customer-managed keys (BYOK), HSM or automated key rotation.

HUMAN APPROVAL

- Per AI employee: asks before everything, before sensitive actions, or not at all. At every level three safeguards stay on: spending cap, AI disclosure to third parties, and logging.
- **Four-eyes:** payments above an amount, deletions, mass sends or permission changes can require two different people, neither the one who assigned the task. Never auto-approved, never in bulk; one rejection blocks. Not yet covered: permission changes made by hand in the dashboard.
- Anything requested by phone waits for a person's approval.

VERIFIABLE ACTIVITY LOG

- Actions by people and AI employees are recorded; the database blocks edits and deletes. Export for any period as PDF, CSV or JSON, **Ed25519-signed with hash-chained rows**.
- An auditor can check an export with a stand-alone script, without Skyseat. It proves the file was not altered after export; by itself it does not prove the database was untouched before.
- Retention is set per company; if the log is limited it never goes below 2 years, and every purge is logged.

IDENTITY AND ACCESS

- **SSO:** OIDC and SAML 2.0 (Okta, Microsoft Entra ID, Google Workspace, ADFS, others). Only DNS-verified email domains; can be made mandatory.
- **SCIM 2.0** provisioning: the provider creates, updates and removes people; groups set the role (never admin).
- **Two-step verification** (TOTP app + 10 single-use recovery codes), enforceable for the whole company; always required for Skyseat operators. Wrong codes are rate-limited per person.
- Roles per company (admin, manager, view-only) plus per-AI-employee permissions, checked by automated sweeps. Login rate limits per email, per IP and in total.
- Sessions: admins set a maximum length (8 h to 30 days; by default 30 days, extended with use) and can sign everyone out. Every sign-in is in the company's Activity log. Admins see and revoke every API key that reaches the company.

BACKUPS AND RECOVERY

- Daily encrypted backup (AES-256) and one before each deploy, stored in a separate data center (Falkenstein): 14 daily + 8 weekly. Weekly automatic restore test.
- RPO up to 24 h. **Disaster drill Oct 31, 2026:** new server rebuilt from the off-site copy alone in 4 min 18 s, all data read; 20–30 min realistic end to end.
- One platform server in one region (no hot standby yet). Status page and monitoring from 4 regions. **No contractual SLA.**

HOW WE USE AI

- Each task sends the model only what it needs. Content read from emails, web pages or files is treated as information, not instructions; sensitive actions need approval. We do not claim prompt injection is fully prevented.
- The platform (not the AI employee) adds an AI disclosure to email, chat, calls and meetings; a company may reword it but not remove it.
- **No training on your data.** Anthropic's commercial API does not train on inputs. The platform is moving to a commercial API or enterprise plan before general availability.

YOUR DATA

- Download everything at any time; on closure all data is deleted after 30 days. DPA (GDPR Art. 28 and a U.S. state-law annex) is drafted and pending lawyer review; breach notice to you within 48 hours of our detecting it, leaving you 72 hours to notify the authority (draft).

WHAT WE DO NOT HAVE YET

- **No SOC 2, ISO 27001, HIPAA, PCI or FedRAMP.** SOC 2 Type I is planned; we have not engaged an auditor and **there is no confirmed date**. A control map and draft policies exist.
- **No external penetration test.** We run our own security reviews (findings fixed, each with a test) and a light outside-in test of our public domains; these are not independent.
- Policies are drafts. No formal HR security program, formal access reviews, DPO or MDM. Continuous WAL backups and hot standby are planned. No `security@` mailbox or `security.txt` yet: report issues from the dashboard (Help and support, reply in 1 business day).
- Some integrations (phone, employees' own email, WhatsApp, Microsoft 365, Google Workspace) are still in test mode. No VPAT or formal accessibility audit.

Skyseat is pre-incorporation; legal entity details will appear on contract documents. This sheet describes controls, not a legal compliance claim. A completed security questionnaire (CAIQ / SIG Lite style) is available on request.