Security

Security brief

What the platform does to protect your workspace, stated from the code and its tests.

Access control

- Hashed API keys per workspace; admin key for operators only.
- Member logins with scrypt passwords, sessions stored as SHA-256, idle and absolute expiry.
- Roles: owner, manager, agent, billing. Default-deny permissions on every route.
- Two-step verification (TOTP) with single-use recovery codes.

Secrets at rest

- Provider tokens and two-step secrets are Fernet-encrypted with a key kept outside the database.
- No API returns a secret, token or ciphertext; logs redact tokens.
- Production refuses to start without its access-control secrets.

Workspace isolation

- Routing matches an exact (account, inbox) pair; anything ambiguous is quarantined.
- Knowledge sources, routing logs and tickets are scoped to the workspace.
- Every sensitive change is written to the audit log without secrets.

Delivery guarantees

- Every webhook is stored before it is processed; failures are retryable.
- A late AI reply can never overwrite a human: ownership is re-checked before every send.
- Provider errors never produce customer-facing text; they hand off with a note.

Data handling

- Conversations live in your Chatwork inbox; the bridge keeps no second copy.
- Order data from Salla and Zid is cached without customer name, address, phone or email.
- Outbound URLs must be https on public addresses (SSRF guard).

Operations

- Self-hosted: Docker Compose, PostgreSQL, Redis, Caddy with automatic TLS.
- Liveness and readiness endpoints; request IDs on every response.
- One gate script: more than 900 tests, database tier included, before any release.

Stated limits

- Email addresses are not verified at signup yet; use code or closed signup mode on public hosts.
- ZATCA: phase-1 QR data only; not a compliance claim.
- The semantic layer has not yet been exercised against a live embeddings provider.