Security Overview

How MeetCrew protects your meetings and your data.

MeetCrew is an AI meeting participant that joins Microsoft Teams, Zoom, Webex, and Google Meet as a named, disclosed teammate and fronts a chat agent you already run. This overview summarizes our security posture for your review. It describes our design at a level intended to answer a security assessment and intentionally omits detail that could aid an attacker. For a security questionnaire, a Data Processing Agreement, or a deeper technical review, contact security@meetcrew.ai.

01 Product & architecture

- MeetCrew is a **meeting-presence layer** that fronts the chat agent you already run (Microsoft Copilot, Copilot Studio, Microsoft Foundry Agent Service, or a custom backend). It is not an agent platform and does not provide the underlying AI reasoning.
- Built on **Microsoft Azure and .NET**. Language inference runs on **in-region Azure OpenAI**; meeting content stays inside the Azure boundary.
- **Data flow:** a meeting-connectivity layer places the named participant into Teams, Zoom, Webex, or Google Meet and streams media in real time under **Zero Data Retention** — the capture vendor stores no recording at any point → MeetCrew's presence layer (avatar, voice, disclosure, memory) runs in your region on Azure → your backend agent generates the substantive responses.

02 Data handling & privacy

- **Processed:** meeting transcripts, structured extractions (topics, decisions, action items, open questions, commitments), and participant platform identities (email and display name from Microsoft Entra).
- **No voiceprints.** No voiceprint or voice-derived biometric identifier is ever created or stored. Cross-meeting recognition uses Microsoft Entra platform identity, never voice.
- **Anonymous speakers** are a first-class concept; named attribution is off by default in v1.
- Customer meeting content is **never used to train any model** — a commitment enforced in our build pipeline (an architectural guardrail test fails the build if a training or fine-tuning code path is introduced against customer content), not by policy alone, and bound onto our content-touching subprocessors by contract. Personal data is not sold or used for advertising.

03 Encryption & key management

- Encryption **in transit** (TLS) and **at rest** on Microsoft Azure.
- **Customer-managed encryption keys (BYOK)** supported for Enterprise: your memory is encrypted at rest under a key held in your own Azure Key Vault or Managed HSM — across all three memory tiers (structured records, raw transcripts, and the semantic search index) — and you can rotate or revoke it. Backend credentials are stored as Azure Key Vault secret references — never as raw secrets.
- Data exports are **double-wrapped**: a tenant key plus your own recipient key.

04 Identity & access

- Authentication through **Microsoft Entra**; access follows your existing identity and conditional-access policies.
- **Tenant-admin authorization** governs which teammates exist, where they may join, and who may configure them.
- Least-privilege access to production; administrative actions are attributed to the acting administrator.

05 Tenancy & isolation

One MeetCrew tenant maps to exactly one Microsoft Entra tenant, with **hard isolation** between customers by construction: structured memory in a per-tenant Azure SQL database, raw artifacts in a per-tenant Blob container, and semantic search in a shared, security-trimmed index that filters every query to the calling tenant. There is no cross-tenant data sharing and no configuration that enables it.

06 Data residency & geo-fencing

- Each deployment is **pinned to a region**, with no cross-region replication of your data.
- v1 regions: **United States (excluding Illinois) and Canada**. The EU, the UK, and Illinois are geo-fenced out; a teammate refuses to join if the organizer's tenant is out of scope.

07 Retention & deletion

- At the **capture layer**, the meeting bot runs **Zero Data Retention** by default — no recording is stored by the vendor at any point — so the only durable copy of your meeting memory is the one in your MeetCrew tenant, governed by your retention policy.
- Retention is configurable per tenant and enforced to the jurisdiction's **statutory ceiling**.
- Expired records are **cryptographically shredded** (the keys protecting them are destroyed). On termination, data is available for export for a limited period, then deleted or shredded, except where retention is required by law.
- **Right to erasure (DSAR).** On an admin-mediated data-subject deletion request, the individual is removed from the **hot, queryable surfaces** — the structured memory and the semantic index your agents and people search. Their name is regenerated out of retained multi-party records (not merely hidden), records that cannot be cleanly regenerated are suppressed, and a durable name-free tombstone prevents any later re-extraction from reintroducing it. The raw source transcript is retained as a bounded, non-queryable residual and is deleted when the request asks for it.

08 Disclosure, provenance & audit

- **Three-layer AI disclosure:** named participant with an AI badge; an on-screen "AI colleague" banner composited into every frame; and an audible self-introduction. A runtime voice opt-out is honored and logged.
- Synthesized speech carries an **audio watermark**; stored artifacts carry **C2PA-style content credentials**.
- Disclosures, opt-outs, and deletions are written to a **hash-chained, tamper-evident audit log** re-verified nightly. Fail-closed: a failed integrity check raises an alert and is never silently overwritten.

09 Subprocessors

We use a small set of vetted subprocessors under written terms, each processing data on our instructions: **Microsoft Azure** (compute, storage, Key Vault), **Azure OpenAI** (in-region inference), **Azure AI Search** (semantic index), and **meeting-connectivity providers** that place the participant into a meeting. Each subprocessor that touches raw meeting content (Recall.ai, Azure OpenAI) is contractually bound by flow-down terms to the same no-training commitment. A current subprocessor list is available on request.

10 Compliance & governance

- MeetCrew is **not yet SOC 2 certified**; a SOC 2 program is on our roadmap, and we support customer security reviews in the meantime.
- The architecture is designed to be defensive against biometric-privacy law (for example, Illinois' BIPA) and data-protection regimes: no voiceprints, disclosure on by default, and geo-fencing where the rules are strictest.
- Two core commitments — **no model training on customer content** and **no voice biometrics** — are enforced in the build pipeline by architectural guardrail tests that fail the build if a violating code path is introduced, not by policy alone.
- A **Data Processing Agreement** is available. For meeting content, MeetCrew acts as your processor and you remain the controller.

11 External interfaces & agent access

- **MCP memory server & API.** MeetCrew exposes structured meeting memory — decisions, open questions, action items, statements, and topics, with platform-identity attribution and timestamps — so a customer's own agent can read it and act on it. Two interfaces are provided: a Model Context Protocol (MCP) memory server and a REST API. Writes are **confined to workflow state** on action items, decisions, and open questions (create an action item, change its status or assignee, resolve an open question, update a decision's status). The surface provides **no delete**, and it **cannot modify transcripts, statements, provenance records, or the audit log**.
- **Authentication.** Access uses **Microsoft Entra** OAuth. The customer grants a service principal in the admin console, and every call is authenticated against that identity. No shared API key substitutes for identity.
- **Scoping & isolation.** Every call is scoped to a single tenant and enforced server-side — a caller cannot read or change another tenant's memory, including by supplying identifiers for records outside its tenant. Reads and write actions are further scoped to a single deployment (one named teammate); cross-deployment access is denied by default and requires an explicit admin opt-in grant.
- **No biometric exposure.** Attribution is derived from Microsoft Entra platform identity. The interface exposes no voiceprints, no speaker audio, and no biometric templates; anonymous speakers are represented as such. Memory served over these interfaces remains within the Azure boundary and is not used to train any model.
- **Auditability.** Reads and write actions over these interfaces are logged and attributable to the authenticating identity.
- **Outbound agent-backend contract.** Separately, during a meeting MeetCrew calls the customer's own backend agent to obtain the spoken reply. This is a distinct outbound interface: the endpoint is customer-configured, the credential is held as an Azure Key Vault secret reference, and the call carries the meeting context required to generate the reply.

12 Reporting a vulnerability

Report suspected vulnerabilities to security@meetcrew.ai. Our responsible-disclosure policy — scope, safe harbor, and what to expect — is published at **meetcrew.ai/security**.

No Compromise AI, LLC · 8 The Green #8797, Dover, DE 19901, United States · security@meetcrew.ai · meetcrew.ai/trust