TEJMA

Privacy Policy

TEJMA LLC · TEJMA DISPATCHER · PRIVACY POLICY

This Privacy Policy explains what information Tejma LLC ("Tejma", "we", "us") collects when you use **Tejma Dispatcher** — our message-ingestion and task-automation service — how we use it, who we share it with, how long we keep it, and the choices you have. **Tejma Dispatcher** exists to process your communications, so please read the "Information We Collect" section carefully: connecting a channel means the messages flowing through it, including messages written by other people, are stored and processed in your account.

Effective August 31, 2026 · Last updated August 31, 2026 · Reno, Nevada

1. Who We Are and What This Covers

Tejma LLC is a Nevada limited liability company based in Reno, Nevada. This policy covers the **Tejma Dispatcher** service — its API, MCP endpoint, dashboard, and background workers. It does not cover third-party services you connect to it (your email provider, Slack, Discord, Anthropic, Stripe, or MCP connector servers you register), which have their own privacy policies.

Our role. For business customers, the organization that connects channels and runs agents is the controller of the content processed through its account; Tejma processes that content on the organization's behalf as described here and, where applicable, under our Data Processing Agreement.

2. Information We Collect

Account information. Sign-in is handled by WorkOS AuthKit. We receive and store your WorkOS user identifier, email address, and organization membership, plus roles your admin assigns. We do not store your password — authentication credentials live with WorkOS.

Messages from connected channels. When you or your organization connects a channel, we receive and store the communications flowing through it: emails (including subject, body, sender and recipient addresses) delivered by our email connectivity provider Nylas; Slack messages delivered through Slack's Events API; and Discord messages streamed through Discord's gateway. For messages that continue a conversation, we also store a bounded plain-text transcript of the earlier messages in that thread, fetched at ingestion, so tasks are understood in context.

Files and attachments. Email attachments and files shared with chat messages (up to 50 MB each) are downloaded and stored as documents on their message, in an object store keyed to your user account.

Content you create. Tasks, agent definitions and skill instructions, extraction and routing skills, MCP connector registrations, and messages your agents send through outbound channels.

Credentials, encrypted. Channel OAuth grants, MCP connector authorization tokens, and your organization's (and any per-user) Anthropic API keys. These are encrypted at rest with AES-256-GCM, decrypted only at the moment of use, and never returned by any API — responses carry only a has-key flag.

Billing information. Your seat's plan history, task-execution usage counts, channel and agent counts against plan caps, overage-credit ledger entries, and Stripe identifiers (customer, subscription, and checkout-session ids). Card numbers are entered on Stripe's hosted pages and never touch our systems.

Logs and diagnostics. Standard server logs (request metadata, timestamps, error traces) and webhook delivery records used to keep ingestion reliable.

3. How We Use Information

We use the information above to run the service you configured: to ingest and store messages, extract the tasks they contain, route each task to your agents, execute approved tasks, send the outbound messages your agents are directed to send, meter task executions against your plan, collect payment, secure the service, and provide support. We also use aggregate, non-content operational data to maintain reliability.

> **We do not sell your data, use it for advertising, or use your content to train AI models.**

Message content is processed only to provide the service to you, by us and by the service providers listed below.

4. AI Processing

Content is sent to Anthropic. To extract tasks from a message, route a task to an agent, or execute a task, the relevant content — the message text, its thread transcript, the task description, and material the agent's skill needs — is sent to Anthropic's Claude API. Every such call is made under your organization's own Anthropic API key (or a per-user key your admin set), so the processing runs under your organization's direct relationship with Anthropic and Anthropic's commercial data terms; there is no shared Tejma key. Under Anthropic's commercial terms, API content is not used to train Anthropic's models absent your opt-in.

Your MCP connectors. When an agent executes with MCP connectors you registered, task-related content may be sent to those connector servers. You choose those servers; review their privacy practices before wiring them to an agent.

5. Service Providers and Recipients

We share information only with the providers needed to run the service, each bound to use it solely to provide their service to us and to you:

- **WorkOS** — authentication, sessions, organization and member management (identity data).
- **Anthropic** — AI processing of message and task content under your organization's own API key.
- **Nylas** — email connectivity: mailbox OAuth, message sync, and attachment retrieval (US region by default; EU region available).
- **Slack** and **Discord** — receiving messages from, and sending messages to, the channels you connect, through their APIs.
- **Stripe** — payment processing: checkout, recurring subscription billing, and dunning (billing contact and payment data).
- **Railway** — cloud hosting: runs our servers and the Postgres database that holds messages, tasks, and account data.
- **Cloudflare** — R2 object storage holding message documents and attachments.

Beyond these, we disclose information only if required by law or legal process (we will notify the organization admin unless legally barred), to protect the rights, safety, or security of Tejma, our customers, or the public, or as part of a merger or acquisition — in which case this policy continues

to apply until you are notified otherwise. A current subprocessor list for business customers is maintained in the Data Processing Agreement.

6. Data Retention and Deletion

Content is kept until you delete it. Messages, their thread transcripts, and their stored documents are retained so your tasks keep working even after the originating email, chat message, or channel connection is gone — disconnecting a channel stops new ingestion but does not delete what was already stored. Content remains until you delete it or your account is closed.

Deletion controls. You can delete stored documents, delete tasks, delete agents and connectors, and disconnect channels or whole channel accounts through the dashboard and API at any time. Deleting a document removes its bytes from the object store before its record. Storage is keyed per user, so a user's stored files can be deleted in bulk when an account is closed.

After account closure. When an organization terminates, we honor export requests for 30 days and then delete the organization's content in the ordinary course of operations. Billing records and ledgers are retained as long as tax and accounting law requires. Encrypted credentials are deleted with the records that hold them; backups age out on a fixed schedule.

7. Security

We protect information with measures matched to what the service holds:

- stored secrets — connector tokens and Anthropic API keys — encrypted at rest with AES-256-GCM and treated as write-only;
- strict per-tenant isolation: every record belongs to one user, and another user's identifiers behave as if they do not exist;
- signature verification on every inbound webhook (Nylas, Slack, and Stripe events are each cryptographically verified before processing);
- OAuth 2.0 bearer-token authentication with JWT verification against WorkOS on every API and MCP request, and TLS in transit;
- admin-only controls for billing, member management, and credential rotation.

No service can promise perfect security. If a breach affects your personal data, we will notify affected organization admins without undue delay and as the law requires.

8. Your Rights and Choices

Access, correction, deletion, portability. You can access and delete most content directly in the service. For anything else — a copy of your data, correction, or full deletion — email us at the address below and we will respond within the time the applicable law sets. If your personal data reached us through a customer's connected channel (for example, you emailed one of our customers), we will refer your request to that customer, who controls that content, and assist them in honoring it.

Region-specific rights. Depending on where you live, you may have statutory rights — under the GDPR or UK GDPR (access, rectification, erasure, restriction, portability, objection, and complaint to a supervisory authority) or the California Consumer Privacy Act (to know, delete, correct, and not be discriminated against; we do not sell or share personal information as the CCPA defines those terms). We honor these rights regardless of which law applies.

Communications. Service and billing notices are part of running the service. We do not send marketing communications through the service's data.

9. Email Provider Account Data

Google. If you connect a Gmail or Google Workspace mailbox, our use of information received from Google APIs adheres to the Google API Services User Data Policy, including its Limited Use requirements: Google user data is used only to provide the message-ingestion and task features you see, is never used for advertising, and is never transferred except to provide those features (including the AI processing described above), for security, or to comply with law.

Microsoft. Microsoft 365 and Outlook mailboxes connected through the service are handled the same way, in line with Microsoft's API terms.

10. Children, International Transfers, and Changes

Children. The service is a business tool and is not directed to children under 16; we do not knowingly collect their data.

International transfers. We operate from the United States and our providers process data primarily in the United States (email connectivity can be pinned to the EU region). Where the GDPR applies to transfers, we rely on appropriate safeguards such as Standard Contractual Clauses, as detailed in the Data Processing Agreement.

Changes to this policy. We will post updates here and change the date at the top; material changes will be announced to organization admins through the service or by email before they take effect.

Contact. Tejma LLC · Reno, Nevada · ai.assistants@tejmallc.com · tejmallc.com

Matej Meciar · Tejma LLC
925-963-4257 · ai.assistants@tejmallc.com · tejmallc.com