Privacy policy

Copy of the Privacy Policy published at https://www.getcargo.ai/privacy.

Last Updated: 29 October 2026

This Privacy Policy describes how GetCargo Inc. ("**Cargo**", "**we**", "**us**") collects, uses, shares, and retains personal data, and the controls available to you. It applies to:

- our websites, including https://www.getcargo.ai and related pages;
- the Cargo cloud product (the "**Services**");
- our hosted Model Context Protocol server at https://mcp.getcargo.io/mcp (the "**MCP server**");
- the Cargo app listed in ChatGPT (the "**ChatGPT app**"), which connects to the MCP server;
- the Cargo app listed in the Slack Marketplace (the "**Slack app**"), including all data Slack sends to Cargo and all data Cargo posts back into Slack; and
- the Google Calendar, Google Sheets, and Google Ads connectors, including all data Cargo receives from Google APIs.

If you are using Cargo under an organization account, your organization is typically the controller of **Customer Content** (the data it puts into Cargo). Cargo is the processor of that data, on the terms of our Data Processing Agreement. Cargo is the controller of account, billing, website, and product-analytics data we collect for our own business.

Contact: legal@getcargo.io. Postal address: GetCargo Inc., 603 Tennessee St, San Francisco, California 94107, United States.

1. Information we collect

1.1 Account and workspace data (Cargo as controller)

When you create an account, join a workspace, or authenticate (including via the ChatGPT app, the Slack app, or another MCP client), we collect:

- identifiers: name, email address, user uuid, workspace uuid and name;
- authentication data: login events, OAuth grants to connected apps (including ChatGPT and Slack), and session tokens;
- billing and plan data: subscription plan, credit balances, invoices, and payment-method metadata processed by our payment provider;
- profile and role data: workspace membership and permissions.

1.2 Customer Content (Cargo as processor)

Depending on how your organization uses the Services, Customer Content may include identifiers, contact details, professional information, CRM and GTM records, communications content, files, knowledge-base documents, workflow configuration, run history, and other data your organization submits or generates.

Categories of data subjects may include your organization's users, employees, customers, prospects, and business contacts.

Cargo does not require special-category or "prohibited" data. Do not submit it unless your contract expressly allows it.

1.3 Website and product analytics

When you visit our websites or use the product, we and our processors may collect IP address, browser and device type, pages viewed, referring URLs, timestamps, and similar usage events. We use Cookiebot for consent, RudderStack and related analytics tools for product usage, and Dub for referral attribution. Advertising cookies load only if you consent.

1.4 Support and communications

If you contact us, we receive the contents of the message, your contact details, and any attachments you send.

2. ChatGPT app and MCP tools

The ChatGPT app is a client of our MCP server. After you authorize it, ChatGPT sends tool inputs to Cargo and receives tool outputs back into the chat. Cargo does not receive the rest of your ChatGPT conversation unless that text is included in a tool argument.

OpenAI is the operator of ChatGPT. Data ChatGPT sends to or receives from Cargo is also processed by OpenAI under OpenAI's terms and privacy policy. Disconnecting the app in ChatGPT or revoking the OAuth grant in Cargo stops further tool calls; it does not delete history OpenAI already holds.

2.1 Tool inputs we receive

Tools may receive, as needed to fulfill the request:

- identity of the signed-in user and bound workspace (from the OAuth token; `whoami` takes no extra input);
- search strings, filters, limits, and timestamps you or ChatGPT supply (`search_actions`, `list_runs`, `query_models`, context search);
- action objects, field values, and record payloads (`get_action_schema`, `autocomplete_action`, `execute_action`, `execute_action_batch`);
- run or batch uuids (`get_run`, `get_batch`);
- model identifiers and SQL `SELECT` text (`describe_model`, `query_models`);
- context paths, globs, and line ranges (`map_context`, `search_context`, `browse_context`, `read_context`).

2.2 Tool outputs we return

Tools return only what is needed for the stated purpose. We do **not** return computed node configs, execution traces, Temporal workflow ids, trace ids, raw run or batch records, or internal object-storage keys.

Tool	Categories returned
`whoami`	User uuid, first name, last name, email; workspace uuid and name; plan name and credit balances
`get_usage`	Credit usage for the last 7 days, grouped by integration, with day timestamps
`search_actions`	Matching actions (name, slug, kind, integration, credit cost, autocomplete hints)
`get_action_schema`	Input and output JSON schemas for an action

Tool	Categories returned
`autocomplete_action`	Options from a connected system (for example Slack channel names, HubSpot object types)
`execute_action`	Outcome; last-node **action output** (Customer Content such as contact or company fields the action produced); run uuid; credits used; or an error message
`execute_action_batch`	Outcome; batch uuid; credits used; a time-limited download URL for the output CSV; or an error message
`get_run`	Same shape as `execute_action` (action output, run uuid, credits). Not nested debug data
`get_batch`	Same shape as `execute_action_batch`. Not the raw batch record
`list_runs`	Run uuid, status, and error message
`list_models` / `describe_model`	Model uuid, slug, name, description, kind, column labels and types
`query_models`	Rows from workspace models. Those rows may include personal data your organization stored (names, emails, titles, and similar)
`map_context` / `search_context` / `browse_context` / `read_context`	Knowledge-base paths, titles, summaries, matching snippets, and file text your organization stored

Action output, SQL rows, autocomplete options, context files, and batch CSVs can contain nested fields (for example a contact object with email, name, and company). Those nested fields are Customer Content, not diagnostic telemetry. Error messages may quote a failed input or a provider error.

3. Slack app

The Slack app lets a Cargo workspace talk inside Slack: mention or DM a Cargo agent, and post messages from plays, tools, or agents into channels and DMs. Installing it (from the Slack Marketplace or by connecting Slack in Cargo) authorizes Cargo to receive Slack event payloads and to post back into that Slack workspace.

Slack is operated by Salesforce. Data Slack sends to or receives from Cargo is also processed by Slack under Slack's terms and privacy policy. Removing the Cargo app in Slack or disconnecting the Slack connector in Cargo stops further event delivery and posts; it does not delete history Slack already holds.

This policy covers all personal data processed through the Slack app. It is not limited to the website.

3.1 Data we receive from Slack

As needed to run the app, Slack sends Cargo:

- Slack team (workspace) ID and the Cargo Slack app ID;
- user IDs, display names, and email addresses of people who install the app or interact with it;
- channel IDs and names for conversations the app is invited to or configured to watch;
- message text, thread timestamps, files, reactions, and related metadata in conversations where the app is mentioned, DM'd, or otherwise invoked;
- OAuth access tokens for the bot (and, if granted, for the installing user).

Cargo does not join every public channel by default. It reads message history in conversations where it is invoked or that a play, tool, or agent is configured to post to.

3.2 Data we send back to Slack

Cargo posts into Slack only what a user or a configured play, tool, or agent asked it to:

- agent replies, streamed into the thread that invoked the agent;
- messages, files, and Block Kit layouts a play or tool is configured to send;
- suggested prompts in Slack's agent container.

Those posts can include Customer Content (names, emails, account fields, research output) that your organization already stored in Cargo.

3.3 How Slack data is used

We use Slack data to authenticate the Slack workspace, bind it to one Cargo workspace, run the agent or play that was invoked, stream or post the result back into Slack, and enforce permissions. We do **not** use Slack data (or any Customer Content) to train foundation models.

4. Google user data

Cargo's Google Calendar, Google Sheets, and Google Ads connectors access a Google account only after a user, or a Google Workspace super admin for domain-wide delegation, signs in with Google and grants access. Revoking Cargo at myaccount.google.com/permissions, removing Cargo's client id from domain-wide delegation in the Google Admin console, or disconnecting the connector in Cargo stops further access.

4.1 Data we access

- **Google Calendar** (read-only): calendar events of the connected user, or of the users a Google Workspace admin selects under domain-wide delegation — titles, times, attendees, descriptions, conferencing links, and status; and, for domain-wide delegation, the domain's user directory (names, email addresses, admin status) so the admin can choose whose calendars to watch.
- **Google Sheets**: the spreadsheets a play, tool, or agent is configured to read from or write to.
- **Google Ads**: the ad accounts, campaigns, and reporting data a play, tool, or agent is configured to read or update.

We also store the OAuth refresh tokens Google issues, encrypted at rest.

4.2 How we use it

We use Google user data only to provide the features the user configured: listing and reading calendar events, waking an agent or play when a meeting is scheduled or changes, reading and writing the selected spreadsheets, and reading or updating the selected ad accounts. Results are stored as Customer Content in that workspace's run history.

We do **not** use Google user data for advertising, sell it, or let humans read it, except with the user's consent, for security or abuse investigations, to comply with law, or in aggregated and anonymized form for internal operations. We do **not** use Google user data to develop, improve, or train generalized AI or machine-learning models. When a workspace runs a Cargo-operated AI feature on Google user data, the relevant data is sent to our model providers (Section 6) only to produce that result, under terms that prohibit training on it.

Cargo's use and transfer to any other app of information received from Google APIs will adhere to the Google API Services User Data Policy, including the Limited Use requirements.

4.3 Retention and deletion

OAuth tokens are deleted within 30 days after the connector is disconnected or the Cargo workspace is deleted, except backups on their ordinary cycle. Google data stored in run history follows the Customer Content retention in Section 7. To request deletion sooner, email legal@getcargo.io.

5. How we use information

We use the information in Sections 1–4 to:

- provide, operate, secure, and support the Services, websites, MCP server, ChatGPT app, and Slack app;
- authenticate you, bind an MCP session or a Slack workspace to one Cargo workspace, and enforce permissions;
- run the actions, queries, agent turns, and context reads you (or an authorized client such as ChatGPT or Slack) request, and return the results to that client;
- meter credits, bill, and prevent fraud or abuse;
- understand product usage and improve the Services (on aggregated or controller-side analytics data, not by training foundation models on Customer Content except as your contract allows);
- communicate with you about the product, security, and — where permitted — marketing;
- comply with law and enforce our terms.

We do not sell personal data. We do not share it for cross-context behavioral advertising except where website advertising cookies run after Cookiebot consent.

6. Recipients

We disclose personal data to:

- **the MCP client you authorized**, including OpenAI when you use the ChatGPT app, which receives the tool inputs and outputs in Section 2;
- **Slack**, when you use the Slack app, which receives the posts in Section 3.2 and which sends Cargo the event payloads in Section 3.1;
- **our subprocessors**, listed in Schedule 3 of the DPA, including Amazon Web Services (hosting), Auth0/Okta (authentication), Temporal (orchestration), Datadog (observability), ClickHouse Cloud (analytics warehouse), and model providers (OpenAI, Anthropic) when you use Cargo-operated AI features;
- **systems your organization connects** (CRM, warehouse, enrichment, Slack, email, and similar). Those parties are your processors, not ours. Tool calls that read or write them send the relevant Customer Content to those systems;
- **payment, support, and analytics vendors** (including RudderStack, Snowflake, Dub, Cookiebot, and our support provider) as needed to operate the business;
- **authorities** when required by law.

Employees access Customer Content only as needed to provide the Services, with confidentiality obligations.

7. Retention

- **Account and workspace data.** Kept for the life of the account. After you delete the account, or after a workspace is closed, this data is deleted within 30 days, except backups on their ordinary cycle and data we

must keep by law (security, billing, and legal claims).

- **Customer Content.** Kept for the term of the contract, then deleted or returned as described in the DPA (typically within 30 days after termination, except backups on their ordinary cycle and data we must keep by law).
- **MCP OAuth grants.** Kept until you or an admin revokes them, or they expire, then deleted within 30 days.
- **Tool call payloads.** Stored in run and batch history as part of providing the Services (so you can inspect what ran), for the same period as Customer Content. MCP responses themselves are not kept as a separate debug archive.
- **Slack OAuth tokens, team IDs, and channel configuration.** Kept until the Slack connector is disconnected or the Cargo workspace is deleted, then deleted within 30 days, except backups on their ordinary cycle.
- **Slack messages, files, and related metadata used as agent input, and the generated reply.** Stored as run history (Customer Content) for the same period as Customer Content. Slack events that do not invoke an agent or a configured play are not kept as a separate archive.
- **Website logs and analytics.** Kept on a rolling basis as needed for security and product analytics, then deleted.
- **Support tickets.** Kept as needed to resolve the request and for a reasonable period afterward, then deleted.

8. User controls

You can:

- review this policy before installing the ChatGPT app or the Slack app;
- disconnect the Cargo app in ChatGPT settings, and revoke the OAuth grant on the MCP server page in Cargo (**Connected apps**);
- remove the Cargo app from a Slack workspace, and disconnect the Slack connector in Cargo;
- update your name and email in account settings;
- ask a workspace admin to remove your membership or delete Customer Content where the product allows it;
- manage website cookies through Cookiebot (banner or cookie declaration);
- request access, correction, deletion, export, or restriction of personal data Cargo holds as controller by emailing legal@getcargo.io. We respond within one month. For Customer Content, we will direct you to your organization (the controller) or assist that organization under the DPA.

California residents also have the rights described in Section 10. We do not sell or share personal information as those terms are defined in the CPRA.

9. Cookies

Cargo uses cookies and similar technologies on the websites:

- **strictly necessary** — authentication, security, and Cookiebot consent;
- **analytics** — to understand how the site and product are used (RudderStack and related tools), if you consent;
- **advertising / attribution** — if you consent (including Dub referral cookies and any advertising partners Cookiebot lists).

You can refuse non-essential cookies in the banner or your browser. Cookiebot's declaration on our site lists the current vendors. Cargo does not use Google's DoubleClick DART cookie.

10. CCPA / CPRA (California)

California consumers may request that we disclose the categories and specific pieces of personal data we have collected as a business, delete that personal data, or correct inaccuracies. We do not sell or share personal data. We will not discriminate against you for exercising these rights. Email legal@getcargo.io. We have 45 days to respond, extendable as the law allows.

Where we process Customer Content as a service provider, those requests should go to your organization. We will not use Customer Content except to provide the Services, for security and debugging, and as otherwise permitted for service providers under the CPRA.

11. GDPR and similar laws

If European Data Protection Laws apply to personal data Cargo processes as controller, you have the right to access, rectification, erasure, restriction, objection, and portability, and the right to lodge a complaint with a supervisory authority. Email legal@getcargo.io. We respond within one month.

Transfers of Customer Content outside the EEA, UK, or Switzerland are handled under the DPA, including Standard Contractual Clauses where required. Subprocessor locations are in DPA Schedule 3.

12. Children

The Services, ChatGPT app, and Slack app are for business users. We do not knowingly collect personal data from children under 16. If you believe we have, contact us and we will delete it.

13. Changes

We may update this policy. We will change the "Last Updated" date above and post the new version at https://www.getcargo.ai/privacy. Material changes that affect the ChatGPT app, the Slack app, or MCP tools will be reflected here before those tools' inputs or outputs change in a way this policy describes.

14. Contact

GetCargo Inc. 603 Tennessee St San Francisco, California 94107 United States

legal@getcargo.io

Related documents: Terms of Service, Data Processing Agreement, Security, Support.