GAL Privacy Policy

Scheduler Systems Ltd
scheduler-systems.com
Confidential

Table of Contents

01 Data Controller

02 What We Collect

03 How We Use Your Data

04 What We Don't Do

05 Analytics

06 Data Storage & Security

07 Data Retention

08 Your Rights

09 International Transfers

10 Customer DPA and Subprocessors

11 Children

12 Changes to This Policy

13 Contact

14 Provider Credentials

SCHEDULER SYSTEMS

Scheduler Systems

Scheduler Systems LLC

scheduler-systems.com

Confidential

Table of Contents

- Data Controller
- What We Collect
- How We Use Your Data
- What We Don't Do
- Analytics
- Data Storage & Security
- Data Retention
- Your Rights
- International Transfers

Effective Date: March 8, 2026 | Last Updated: March 14, 2026

Scheduler Systems Ltd. ("we", "us", "our") operates GAL (gal-run), a developer platform for managing AI agent configurations. This Privacy Policy explains what data we collect, how we use it, and your rights.

---

01 Data Controller

Scheduler Systems Ltd.

Company Number: 517139382 (Israel)

Menachem Begin 150, Tel Aviv, Israel

Email: privacy@scheduler-systems.com

Website: https://gal.run

02 What We Collect

Account Information

When you create an account, we collect:

- Email address
- Name (optional)
- Authentication tokens from GitHub/Google if you use SSO

Usage Data

We collect data about how you use GAL to improve the service:

- Feature usage (clicks, time spent, navigation patterns)
- Device and browser information
- Error logs for debugging
- API request metadata (endpoints called, response times)

CLI Tool

When you use the gal CLI:

- We process authentication tokens to verify your identity
- We collect minimal diagnostic events (command run, success/failure)
- **We do not collect your source code or keystrokes** unless you explicitly send content to our servers

GitHub App

When you install the GAL GitHub App:

- We access only repositories you explicitly authorize
- We read repository metadata needed for integration
- We do not store your source code beyond what's needed for active processing

IDE Extensions

Our VS Code and other IDE extensions:

- Authenticate with your GAL account
- Send minimal usage telemetry (feature used, errors)
- **Do not read or transmit your code** unless you explicitly trigger a GAL feature

Billing Information

When you subscribe to a paid GAL plan, we collect and process:

- Email address and organization name for billing purposes
- Subscription plan tier and status
- Coupon or promotional code usage
- Design partner status (if applicable)

Payment processing is handled by Stripe, Inc. Stripe collects and processes payment card information according to their Privacy Policy. We do not store your full payment card details.

---

03 How We Use Your Data

We use your data to:

- Provide and operate GAL
- Authenticate you across our services
- Improve features based on usage patterns
- Debug issues and fix bugs
- Send important service communications
- Prevent abuse and ensure security

---

04 What We Don't Do

- **We do not sell your data**
- **We do not use your content to train AI models** without your explicit consent
- **We do not share your data with advertisers**
- **We do not read your source code** unless you explicitly send it to our services

05 Analytics

We use PostHog for product analytics. This helps us understand how people use GAL so we can improve it. Analytics data includes usage patterns, device info, and error logs — but not your code or sensitive content.

You can opt out of analytics in your account settings.

06 Data Storage & Security

Your data is stored securely using modern encryption and access controls:

- Data encrypted in transit (TLS) and at rest
- Access limited to authorized personnel
- Regular security audits
- Infrastructure hosted on secure cloud providers

07 Data Retention

- **Account data:** Kept while your account is active, deleted within 30 days of account deletion
- **Usage analytics:** Retained for 24 months
- **Diagnostic logs:** Retained for 12 months
- **Processed content, work items, and session metadata: Retained as needed to operate the Services, maintain work-item history, support security and audit records, respond to deletion requests, and comply with legal obligations.**
- **Billing and subscription data:** Retained for the duration of your account plus 12 months for tax and legal compliance
- **Coupon and promotional data:** Retained for 24 months

08 Your Rights

Depending on your location, you may have rights to:

- **Access** your personal data
- **Correct** inaccurate data
- **Delete** your data
- **Export** your data in a portable format
- **Object** to certain processing
- **Withdraw consent** where processing is based on consent

To exercise these rights, contact us at privacy@scheduler-systems.com.

09 International Transfers

Your data may be processed in Israel, the European Union, the United States, and other jurisdictions where our subprocessors operate. For GAL, production user data is stored in Google/Firebase EU infrastructure. Where personal data is transferred outside the EEA, UK, Switzerland, or another protected jurisdiction, we use lawful safeguards such as adequacy decisions or Standard Contractual Clauses where required.

10 Customer DPA and Subprocessors

If you use GAL as part of a business or organization, our GAL Customer Data Processing Addendum is available at https://scheduler-systems.com/legal/en/gal-dpa.pdf. It describes how we process Customer Personal Data on behalf of business customers and includes the current subprocessor list.

Current GAL subprocessors and service providers include Google Cloud / Firebase for hosting and storage; GitHub for repository and workflow execution; customer-authorized AI providers such as OpenAI, Anthropic, and Google; Intercom for support chat; PostHog for product analytics; Stripe for billing and subscription management; and Brevo for transactional email delivery.

We will provide at least thirty (30) days notice before authorizing a new GAL subprocessor whose engagement materially expands the scope of Customer Personal Data processing.

11 Children

GAL is not intended for users under 16. We don't knowingly collect data from children.

12 Changes to This Policy

We may update this policy. For material changes, we'll notify you via email or through the platform. Continued use after changes means you accept the updated policy.

13 Contact

Questions about privacy? Contact us:

Scheduler Systems Ltd.

Company Number: 517139382

Privacy: privacy@scheduler-systems.com

Website: https://gal.run

14 Provider Credentials

When you connect third-party AI model providers (such as Anthropic, OpenAI, and Google) to GAL, we process the API keys you provide for those providers. This section supplements the general disclosures above.

What we collect. Provider API keys you enter into the Services; a provider account identifier where the provider exposes one; and metadata about which provider credentials exist for your account. Key contents are never recorded in logs or analytics.

Why we process it. To execute background agents you dispatch through the Services on your behalf. The lawful basis is performance of the contract you enter into with us by using the Services (GDPR Article 6(1)(b)). We refuse to dispatch a background agent if a required provider API key is missing, consistent with Section 14 (Provider Credentials - API Keys Required) of our Terms of Service.

How it flows. Credentials are encrypted at rest in Google Firestore. During agent execution, credentials are transmitted to our runner subprocessor (GitHub Actions) and injected only as environment variables for the duration of a single session. Credentials are never written to logs, analytics, or any persistent store outside our encrypted credential repository. We do not use provider credentials for any purpose other than executing the agents you dispatch.

Who receives it. GitHub, Inc. receives the credential as an environment variable for each session; the third-party model provider itself receives the API calls you authorize by storing its key with us; no other subprocessor receives the credential itself.

Retention. Stored until you delete or rotate them via the Services dashboard or CLI. You may revoke a stored credential at any time. Upon account deletion, credentials are removed from our active systems within thirty (30) days; copies in operational backups are purged per our backup retention schedule, typically about ninety (90) days.

Your rights. You may view which credentials are stored against your account, delete any stored credential at any time, and, where offered at your plan tier, request an audit log of when your credentials were used. These rights are in addition to those listed in the Your Rights section above.