

V1.0

Status	Effective April 28, 2026
Effective	April 28, 2026
Canonical source	https://secapi.ai/trust

Service Level Agreement (SLA)

Version 1.0 — Effective April 28, 2026. This SLA is incorporated by reference into the underlying agreement between Provider and Customer and is binding on the Parties as of the effective date. A downloadable PDF is provided for procurement and legal review at </trust/sla.pdf>.

This Service Level Agreement ("SLA") sets out the uptime targets, support response targets, and service-credit remedies that apply to Customer's use of the secapi.ai Services ("Services"). It is incorporated into and forms part of the agreement between secapi.ai ("Provider") and Customer ("Customer"). Capitalized terms used but not defined in this SLA have the meanings given to them in the underlying agreement or the Data Processing Agreement.

1. Tier definitions

This SLA applies to Customer's plan tier as identified in the dashboard or order form:

- **Commercial** — paid contract plan with negotiated terms (annual or multi-year).
- **Team** — self-serve paid plan for teams.
- **Personal** — self-serve paid plan for individuals.
- **Free / PAYG** — free tier and pay-as-you-go consumption without a paid commitment.

The SLA targets and remedies that apply depend on the Customer's tier as set out in Section 2.

2. Service Level targets

2.1 Uptime targets

Tier	Monthly uptime target	Service-credit remedy
Commercial	99.9%	Service credits per Section 4
Team	99.5%	None (support-response targets per Section 3)
Personal	99.5%	None (support-response targets per Section 3)
Free / PAYG	Best-effort	None

2.2 What "uptime" means

"Uptime" means the percentage of time during a calendar month that the Services are available to Customer, calculated as:

$$\text{Uptime \%} = ((\text{Total Minutes in Month} - \text{Downtime Minutes}) / \text{Total Minutes in Month}) \times 100$$

Downtime means a period during which the Production API (`api.secapi.ai`) returns a sustained 5xx error rate above 5% across all healthy regions for at least five (5) consecutive minutes, as measured by Provider's monitoring infrastructure (BetterStack and internal probes; see <https://secapi.ai/status>).

2.3 What is excluded from "downtime"

The following events are NOT counted as Downtime for purposes of this SLA:

- Scheduled maintenance windows announced at least 48 hours in advance via the status page or the Customer's dashboard.
- Emergency maintenance reasonably required to address a security vulnerability or critical bug, where prior notice is not practical.
- Failures attributable to Customer's actions or omissions (e.g., misconfigured webhooks, exhausted rate limits, invalid API keys).
- Failures of third-party services that are outside Provider's reasonable control (e.g., Customer's network, Customer's chosen LLM provider for optional enrichment, the SEC EDGAR upstream when Customer requests live pass-through).
- Force majeure events (e.g., natural disasters, war, government action, internet backbone outages outside Provider's control).
- Beta, preview, or alpha features explicitly labeled as such.

3. Support response targets

The following targets describe the time within which Provider will provide a substantive first response to a support request submitted via the dashboard support widget or `support@secapi.ai` . They are targets, not absolute guarantees.

Tier	Critical (P0)	High (P1)	Normal (P2)	Low (P3)
Commercial	1 business hour	4 business hours	1 business day	3 business days
Team	4 business hours	1 business day	2 business days	5 business days
Personal	1 business day	2 business days	5 business days	Best effort
Free / PAYG	Best effort	Best effort	Best effort	Best effort

Severity definitions:

- **Critical (P0)** – Production Services completely unavailable to Customer; no workaround.
- **High (P1)** – Major feature or workflow severely impaired; significant business impact.
- **Normal (P2)** – Defect or limitation with workaround; moderate business impact.
- **Low (P3)** – Cosmetic issue, documentation question, or feature request.

"**Business hours**" means 9:00 a.m. to 6:00 p.m. United States Pacific Time, Monday through Friday, excluding US federal holidays.

Critical (P0) responses are provided 24x7 for Commercial-tier customers.

Critical bug acknowledgement target across all paid tiers: 24 hours.

4. Service credits (Commercial tier only)

4.1 Eligibility

Service credits are available only to Commercial-tier customers and only on a properly submitted claim per Section 5.

4.2 Credit schedule

Monthly uptime	Service credit (% of monthly fee)
< 99.9% and ≥ 99.0%	10%
< 99.0% and ≥ 95.0%	25%
< 95.0%	50%

"**Monthly fee**" means the prorated monthly portion of Customer's recurring fees for the Services for the calendar month in which the Downtime occurred, excluding overage usage charges, professional services, and one-time fees.

4.3 Form of credit

Service credits are issued as a credit against future invoices and may not be redeemed for cash. Credits do not roll over beyond the contract term and have no monetary value upon termination of the underlying agreement.

4.4 Maximum monthly credit

Total service credits awarded for any single calendar month will not exceed 50% of the monthly fee for that month.

4.5 Sole and exclusive remedy

Service credits are Customer's sole and exclusive remedy, and Provider's sole and exclusive liability, for any failure to meet the uptime targets in Section 2.1. This Section 4.5 does not limit either Party's liability for breach of confidentiality, indemnification obligations, or any other matter for which limitation of liability is excluded under the underlying agreement or applicable law.

5. Service-credit claim procedure

5.1 Submission window

Customer must submit a written service-credit claim to support@secapi.ai within thirty (30) days of the end of the calendar month in which the Downtime occurred. Claims submitted after this window are forfeit.

5.2 Claim contents

The claim must include:

1. The Customer's account email or organization identifier.
2. The dates and approximate times during which the Customer believes Downtime occurred.
3. The Customer's calculation of the Downtime percentage and resulting credit, with supporting evidence (e.g., logs from the Customer's monitoring system, screenshots of error responses with timestamps and request IDs).
4. The impact to the Customer's business (qualitative is acceptable).

5.3 Provider review

Provider will review the claim and respond within ten (10) business days. Provider's monitoring data (as displayed on <https://secapi.ai/status> and Provider's internal observability systems) is the authoritative source for Downtime measurement under this SLA. If Provider's records and Customer's records differ, the Parties will work in good faith to reconcile.

5.4 Escalation

If a claim is denied and the Customer disagrees, the Customer may escalate to the named Provider account executive or to support@secapi.ai with subject line "SLA escalation." Escalations are reviewed by senior operations leadership.

6. Webhook delivery SLA

The webhook delivery SLA (99.9% delivery success, < 5 second first-attempt latency, 72-hour retry window with HMAC-signed payloads) is documented separately at <https://docs.secapi.ai/webhook-sla> and applies to all paid tiers. The webhook delivery SLA is independent of the API uptime SLA in this document.

7. Status communication

7.1 Live status

The current state of the Services and the active-incident history are published at <https://secapi.ai/status> (BetterStack-backed live dashboard). The status page also exposes an RSS feed at <https://secapi.ai/status.rss> for programmatic monitoring.

7.2 Incident communication

For incidents affecting Customer Services, Provider will:

1. Post an initial acknowledgement on the status page within 15 minutes of detection for PO/P1 incidents.
2. Post substantive progress updates on the status page no less frequently than every 30 minutes during a PO/P1 incident.
3. Post a resolution notice on the status page once Service is restored.
4. Publish a post-incident review for PO incidents within ten (10) business days of resolution.

7.3 Direct notifications

Commercial-tier customers may designate a notification email address in the dashboard to receive direct status notifications for incidents affecting their tier.

8. Updates to this SLA

Provider may update this SLA from time to time. Material changes that reduce SLA targets or remedies will be communicated to Commercial-tier customers in writing at least 60 days in advance. The current version of this SLA is published at <https://secapi.ai/trust#sla>.

9. Definitions

- **"Production API"** means the production endpoints served from api.secapi.ai (and any successor domain Provider designates). It excludes sandbox, staging, beta, preview, alpha, and explicitly experimental endpoints.
- **"Region"** means a geographic deployment of Provider's infrastructure (e.g., US-East, US-West).
- **"Healthy region"** means a region where Provider's infrastructure is available and serving traffic. A failure of one region while the global service remains available (via routing or failover) does not by itself constitute Downtime.

10. Contact

- **Service-credit claims:** support@secapi.ai
- **Incident escalation (Commercial):** named account executive or support@secapi.ai with subject "SLA escalation"
- **Live status:** <https://secapi.ai/status>

