Service-Specific Terms

Effective 22 Nov 2026 Version 1.0 For everyone

The short version

- The Coding Plan has no token cap and doesn't use Credits, but it isn't unlimited speed: above fair-use budgets, requests slow down.
- It is for one person's own interactive use; sharing or bulk use can be throttled or stopped.
- It isn't refundable once activated, and a payment dispute suspends its keys.
- NeutrinoPass and Rewards are free-tier boosts with no cash value.
- Router caching and fallbacks are opt-in settings you control.
- MCP servers you add are third-party services; you choose what they can do.

This summary helps you find your way. The full text below is what applies.

These terms add to the Terms of Service for specific products and features. If they conflict with the Terms for that product or feature, these terms apply.

1. Coding Plan (DevPass)

1.1 **What it is.** The Coding Plan is a flat monthly subscription that gives one person access, through dedicated keys, to a set of Models for software development. Coding Plan keys never draw on your Credits. The Models, prices, seat limits, concurrency limits and fair-use budgets are shown on the plan page and in our documentation, and your current usage is shown in your dashboard.

1.2 **What "unlimited" means.** Where the Coding Plan is described as "unlimited" or "unlimited tokens", that means **your requests are not capped by a token allowance and do not use Credits. It does not mean unlimited speed, concurrency or priority.** The Coding Plan runs

on shared capacity that we operate or reserve, priced so that it stays affordable for every subscriber. To keep it usable for everyone:

- **Fair-use budgets.** Each plan has daily and weekly fair-use budgets. Above a budget, your requests keep working but may be admitted more slowly until the budget resets.
- **Concurrency.** Each plan has a limit on simultaneous requests, and larger Models may count as more than one.
- **Low-interactivity mode.** If usage looks far beyond one person coding (for example, sustained unattended or parallel traffic), requests may be slowed for about 24 hours. This is not a suspension.

Slowing requests within these published limits is how the plan is designed to work. **It is not a failure to provide the service, and it does not entitle you to a refund or partial refund.**

1.3 **One person, personal, interactive use.** The Coding Plan is for one individual's own, personal, non-commercial interactive coding, as you confirm when you subscribe. Business or commercial use needs Credits or a separate agreement. Sharing keys, pooling use, automated reselling, or traffic that indicates several people or unattended bulk use may lead to throttling, suspension or termination.

1.4 **Value received.** Usage on the Coding Plan can be worth many times the subscription price at the Models' public per-token prices. That is the purpose of the plan, and it is not a promise that any level of usage will be available at any particular speed. You agree that the service is delivered each time a request is processed, and that the usage records in your dashboard and our systems (requests, tokens, Models, timestamps, and the IP addresses and networks used) are a record of that delivery.

1.5 **Fair-use review.** To enforce this, we keep short samples of recent requests, prompt fingerprints, working-directory paths sent by coding tools, request timing, and the IP addresses and networks used, as described in the Privacy Policy. Our systems slow requests automatically above fair-use budgets; suspensions and terminations are decided by a person, and you can appeal under Terms Section 15.5.

1.6 **Seats, pricing and changes.** Seats may be limited and offered through a waitlist. Where a launch price is promised to keep while you stay subscribed, it lasts until you cancel or your subscription lapses. We may change the Models, budgets or limits of the Coding Plan with notice on the plan page; if a change materially reduces what your plan offers, you may cancel before your next renewal.

1.7 **Refunds. The Coding Plan is not refundable once activated**, except where the Refund & Billing Policy or the law requires. Where a refund or withdrawal is calculated by deducting

what you used, the value of your usage is calculated at the public per-token prices of the Models you used, up to the price you paid.

1.8 **Disputes.** If you dispute a Coding Plan payment after using the plan, we will suspend the plan's keys while the dispute is open and respond to the dispute with the records in section 1.4, the plan terms and fair-use information shown when you subscribed, your acceptance of them, and our support correspondence with you, as set out in Refund & Billing Policy section 7. Please contact support@electronhub.ai first.

2. Neutrinos, NeutrinoPass and Rewards

2.1 **Neutrinos.** Neutrinos are a free daily allowance for Models marked ":free". They reset daily at 21:00 UTC, do not carry over, and have no monetary value. Neutrinos are deducted when a request is accepted, and are not returned if the request later fails or is cancelled. Some free Models require a verified account.

2.2 **Free queue.** Free Models share limited capacity. Requests may be queued, slowed or refused, and paid plans may be served first.

2.3 **Rewards.** You can earn extra Neutrinos by viewing third-party ad or link pages. Rewards are checked for fraud, including device and network checks, and may be refused or revoked for automated, duplicated, VPN-masked or multiple-account activity. Repeated attempts to claim Rewards from several accounts can lead to account suspension.

2.4 **NeutrinoPass.** NeutrinoPass increases your daily Neutrinos and removes the need to view ads while it is active. If the pass is refunded, disputed, paused or unpaid, its benefits end immediately and that day's Neutrinos may be reduced to the free amount.

3. Router and Auto

3.1 **Routing.** When you choose Auto or a Router policy, we pick a Model and route for each request based on the request, your policy and current availability, and charge a small router fee shown in the Service in addition to the chosen Model's usage.

3.2 **Fallbacks.** If you enable fallbacks or provide a list of Models, a failed request may be retried on another Model or provider you allowed.

3.3 **Response cache.** If you turn on response caching, we store responses for the period you choose (up to 30 days) and may answer matching requests from the cache, without calling a Model. Cached answers may be out of date. Cache hits are not charged for Model usage.

4. MCP servers and tools

4.1 **Third-party servers.** Model Context Protocol servers you add are run by third parties or by you. They may read, change or send data, or take actions, when a Model calls them. You decide which servers to add and what they may do, and you are responsible for their use.

4.2 **Credentials.** Connection details and authorisation tokens for MCP servers pass through our gateway. We store OAuth tokens encrypted and use header secrets only for the request that needs them. Revoke access at the server if you no longer use it.

4.3 **Agent tools.** Agents can browse the web, run code in your browser, and call tools. Review what an agent proposes to do before you let it act on accounts or data that matter.

5. Beta and preview features

Features or Models labelled beta, preview, experimental or similar are provided for evaluation. They may change or be withdrawn at any time, may be less reliable, and may have stricter limits.