

What a sandbox provider listing shows

Every listing on this site records the specs a provider publishes on its own docs or pricing page: isolation, cold start, persistence, GPU, pricing and more. Each row links the cited source.

Updated October 20, 2026

Every listing on this site records the facts a sandbox provider publishes on its own documentation or pricing page. Each row carries a source URL that points to the page the fact came from. The listing is a snapshot, not an endorsement.

The name and description identify the provider and what it does. The description is a short summary written for this index. Categories tag the provider by use case: code execution, dev environments, browser sessions, computer use, evals and RL, or open source.

The isolation attribute tells you how the provider keeps one sandbox from interfering with another. The options in this index are MicroVM, container, gVisor, VM, WebAssembly, process, and mixed or selectable. MicroVM is the most common type among providers here. Mixed means the provider offers more than one type or lets you choose.

Cold start is the time a new sandbox takes before it can run code. It is shown in milliseconds. Only a few providers in the index publish this number. Freestyle lists 65 milliseconds, Daytona 90 milliseconds and E2B 150 milliseconds. An empty field does not mean the provider has no cold start cost. It means the number is not published.

GPU support is a simple yes or no. Most rows leave it empty.

Persistence describes how long a sandbox keeps its state. Ephemeral sandboxes start empty every time. Pause and resume lets you stop a session and keep the state for later. Snapshots save the full state on demand. Persistent sandboxes keep their state. Each provider defines these terms in its own docs. The pricing model attribute tells you how the provider charges. Usage-based billing charges for what you use. Plan-based billing sells plans. Hybrid providers combine the two. The free tier attribute shows whether the provider has a free option: free credits, an always-free tier, or an open source licence.

Open source and self-host flags tell you whether you can inspect and run the sandbox software yourself. Several providers in this index publish their sandbox runtime under the MIT, Apache-2.0 or GPL-3.0 licence. The self-host flag means you can deploy the sandbox on your own infrastructure, sometimes through BYOC or VPC deployment.

SDK languages list the client libraries the provider publishes. TypeScript and Python are the most common among providers here. A few also support Go, Ruby or Java. If your agent uses a different language, check whether the provider has an HTTP API or a CLI.

Compute and memory rates are shown as dollars per 10,000 hours. The listing multiplies the per-hour rate by 10,000 so that numbers like 0.0504 keep their digits. 504 on the listing means 0.0504 dollars per vCPU-hour. The rate field is empty when the provider meters a different unit, such as active CPU time or GB-RAM-seconds.

The documentation link points to the provider's own docs or developer guide. It is the most

useful source for questions the listing does not answer.

Frequently asked questions

Why are some fields empty on a listing?

A blank field means the vendor does not publish that spec on its docs or pricing page. It never means the feature is absent. Ask the vendor for the missing number.

What do the compute and memory rates mean?

They show the published list rate in dollars per 10,000 hours. 504 means 0.0504 dollars per vCPU-hour. Some providers bill a different unit, like per active CPU second or per GB-RAM-second, so a missing rate is a metering difference not a gap in checking.

How often are listings updated?

Sources are checked monthly against each provider's documentation and pricing page. The source URL on each listing lets you verify the facts yourself.

Sources

[Vendor documentation and pricing pages, as described on our About page](https://agentsandboxes.fyi/about) — <https://agentsandboxes.fyi/about>