

How to compare sandbox providers

Compare sandbox providers by isolation, cold start, persistence, GPU support, pricing model and free tier using a standard checklist and the compare tool on this site.

Updated October 20, 2026

Start with what kind of sandbox you need. The index organises providers into categories: code execution, dev environments, browser sessions, computer use, evals and RL, and open source. Each category groups providers that share a job. A provider can appear in more than one category when it serves multiple use cases.

Pick two or three providers from the same category. Open each listing and look at the attributes. The isolation column tells you how the provider keeps sandboxes apart. MicroVM is the most common type in this index. E2B and Vercel Sandbox describe their sandboxes as Firecracker microVMs. The index also records containers, full VMs, gVisor, WebAssembly and process-level isolation.

Check the cold start time when speed matters. A few providers publish their boot times: Freestyle lists 65 milliseconds, Daytona 90 milliseconds and E2B 150 milliseconds. Most providers do not publish a cold start figure. That is not a sign of a slow product, only that the vendor does not state the number on its pricing or docs page.

Decide whether you need GPU support. Only a few providers in the index record GPU support. If you run inference, vision models or rendering inside a sandbox, filter for providers with GPU support on, then check the provider page.

Choose your persistence model. Ephemeral sandboxes start fresh every time. Pause and resume lets you stop a session and come back. Snapshots let you restore saved state later. Persistent sandboxes keep their state. Each provider defines these terms in its own docs. Billing differs for each choice, so check how the provider charges for stopped or idle sandboxes.

Understand the pricing model. Usage-based providers bill for the compute you use. Plan-based providers sell plans. Hybrid providers combine a plan with usage. Free tiers exist on many providers in this index, from free credits to always-free tiers. The sandbox cost estimator on this site can help you project a monthly bill from your run count, average run length and the provider's published rates.

Look at the SDK languages. Most providers in the index publish a TypeScript or Python SDK. A few offer Go, Ruby, Java or a CLI tool. Check that the provider supports the language your agent is written in.

Check whether the provider is open source and whether you can run it yourself. Several providers in the index are MIT, Apache-2.0 or GPL-3.0 licensed and can be self-hosted. That matters if you want to keep code on your own infrastructure or need to customise the sandbox runtime.

Use the compare tool on this site to line up two or three providers side by side. The table shows isolation, cold start, GPU support, persistence, pricing, free tier, compute and memory rates, and open source status. An empty cell means the vendor does not publish that number. Rates are shown as dollars per 10,000 hours so that small per-hour rates keep their digits.

After you narrow the list, visit each vendor's documentation link from the listing and read the details yourself. Specs and pricing change. The listing records what the vendor stated on the date checked.

Frequently asked questions

How many providers should I compare?

Two or three is enough to see the differences that matter for your use case. You can compare up to three at a time in the compare tool on this site.

Does a paid badge mean the provider is better?

No. Verified, Featured and Sponsored placements never change the recorded facts on a listing. Compare the spec rows, not the badge.

What if a spec is empty on a listing?

An empty field means the vendor does not publish that number. It does not mean the provider lacks the feature. Ask the vendor directly.

Sources

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