

Common mistakes and questions about sandbox providers

Six mistakes that teams make when picking a cloud sandbox provider and the questions to ask before you commit, based on specs recorded in this index.

Updated October 20, 2026

First mistake: comparing only the per-hour rate. The compute and memory rates on each listing show the published per-hour price in a common unit. But providers meter differently. Some bill per active CPU second, some per 10 milliseconds of CPU time, and some per GB-RAM-second. A low per-hour rate means little if your workload matches a different metering unit. The sandbox cost estimator on this site lets you enter your own run count, run length and sandbox size, then shows the estimated total across the rates you set.

Second mistake: ignoring the persistence model. Ephemeral sandboxes start from scratch in each new session. If your agent builds state over time, look for pause and resume, snapshot or persistent sandboxes. The persistence attribute on each listing tells you which model the provider uses. The right choice depends on whether your agent keeps working memory or starts fresh on every task.

Third mistake: assuming every provider works for every agent language. Most providers in this index publish a TypeScript or Python SDK. A smaller set supports Go, Ruby, Java or a CLI. If your agent is written in a language the provider does not support, check whether the provider has an HTTP API. The SDK languages attribute on each listing lists the supported clients.

Fourth mistake: trusting the per-hour rate without checking idle billing. Some providers bill only when a sandbox actively uses CPU. Others bill for as long as a sandbox exists. The difference can change the real cost. The source link on each listing is the place to start when you check how the provider bills idle sandboxes.

Fifth mistake: ignoring the self-host option. Several providers in the index let you run the sandbox runtime on your own infrastructure. Open source providers publish their code under MIT, Apache-2.0 or GPL-3.0 licences. If you need to keep sandbox data inside a specific region or behind a firewall, filter for providers with open source and self-host flags set to yes.

Sixth mistake: picking a provider before you understand isolation requirements. If your sandbox runs untrusted code from users, isolation type matters. The options in this index are MicroVM, container, gVisor, VM, WebAssembly and process. The isolation attribute on each listing states the type the provider describes. Ask each provider about its security practices if isolation is critical to your security model.

Questions to ask a provider before you commit: what is the cold start time for a new sandbox, does it bill idle time, can you choose a region for your sandbox servers, what happens to data when a sandbox stops, how long are snapshots kept, what is the support response time, and does the provider publish a roadmap of breaking changes. The documentation link on the listing is the first place to look. The index does not record every answer.

Frequently asked questions

Can I switch providers after I build on one?

Switching may be easier when your agent talks to the sandbox through an SDK or HTTP API in a language that several providers support. Check the SDK languages attribute on each listing.

Does a free tier mean I can run production on it?

Free tiers usually come with limits on runtime, memory or concurrent sessions. The listing records the free tier type but not the specific limits. Check the provider's pricing page before you rely on a free tier for production work.

Should I pick the cheapest provider?

The cheapest per-hour rate may not be the cheapest total. Free credits, included usage, storage fees, data egress charges and idle billing all affect the real bill. Use the sandbox cost estimator to project your workload at the rates you enter from each provider's page.

Sources

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