

MCP SERVER

NO CODE

CLOUD HOSTED

E2B MCP for AI Agents

Run python and shell commands in secure cloud sandboxes.

E2B provides secure cloud sandboxes for AI code execution. It lets your agent run Python, JavaScript, and shell commands in isolated Firecracker microVMs with 150ms cold starts. It's built for speed and safety.

A+ Quality Score 100/100

sandboxing

code-execution

micro-vms

firecracker

ai-agents

isolated-environment



The connectivity layer between AI and the world's software.

Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

Your AI Connections Run Through Vinkius Cloud

The world's largest
managed MCP catalog

Vinkius is the connectivity layer where AI connects to the software your business already runs. We handle the hosting, the security, the credentials, the uptime — you get agents that actually do things.

We operate the world's largest managed MCP catalog. Major SaaS platforms, CRMs, databases, and cloud providers — running, monitored, production-ready. This MCP server is hosted and maintained by the Vinkius Cloud for AI Agents.

The agent doesn't manage credentials, doesn't manage uptime, doesn't manage security. Vinkius does.

— Architecture principle

Four Pillars of the Vinkius Runtime

01 — Security by design

Credentials stay encrypted at rest via AES-256. The AI agent never touches raw keys — they're injected into a sandboxed V8 isolate at runtime. Actions are logged, and connections have an emergency kill switch.

03 — Deterministic observability

Eight immutable metrics per endpoint: request volume, p95 latency, error rate, active connections, cost attribution. A live payload feed logs every tool call with mutation detection.

02 — Built on MCP Fusion

This MCP server was built with **MCP Fusion**, the open-source framework (Apache 2.0) that powers the entire Vinkius catalog. Schema-as-firewall strips undeclared fields, compiled PII redaction runs at zero overhead, and cryptographic lockfiles produce git-diffable audit trails.

04 — Autonomous operations

Servers are deployed, monitored, and patched autonomously. New capabilities and security patches ship weekly. Zero-downtime deployments ensure continuous availability across all managed MCP servers.

AES-256

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

Sandboxed execution

One Token. Instant Access.

Every MCP server on Vinkius is accessed through a **Connection Token**. Tokens are generated in the cloud dashboard and produce a unique MCP endpoint URL. Paste this URL into any MCP-compatible client — no SDK required.

A single token can serve **multiple AI clients simultaneously**, or you can issue separate tokens per client for granular access control. Each token tracks its own request count, last activity timestamp, and can be individually enabled or revoked.

MCP ENDPOINT`https://edge.vinkius.com/{token}/mcp`

Claude



Cursor



VS Code



Windsurf



Grok



Gemini

Security Is the Architecture

Security in Vinkius is not a feature — it's the foundation of the runtime. The gateway enforces multiple independent protection layers between AI agents and third-party APIs.

01 — Ed25519 PKI Vault

Every workspace has an Ed25519 Master Key. Session keys are generated ephemerally (24h TTL) and signed by the Master Key. Credentials never leave the vault boundary.

02 — V8 Isolate Sandboxing

Tool code runs inside isolated-vm V8 isolates with 64 MB memory caps and per-request timeouts. No filesystem access, no network access except through the SSRF-guarded fetch bridge.

03 — SSRF Guard

All outbound HTTP requests are DNS-resolved and validated before execution. Private IP ranges (10.x, 172.16-31.x, 192.168.x, AWS metadata 169.254.x) are blocked at the network layer.

05 — Cryptographic Audit Trail

Every request is signed into a SHA-256 hash chain with Ed25519 signatures. Events form a tamper-proof, SIEM-exportable forensic record.

04 — DLP & PII Redaction

A ResponseGuard pipeline intercepts every tool response. Configurable redaction patterns strip sensitive fields (emails, SSNs, card numbers) before data reaches the AI agent.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot is used outside Vinkius infrastructure, the server is quarantined instantly.

Emergency Kill Switch

EU AI Act Art. 14(1)
Compliant

The kill switch is an **emergency halt** mechanism — not a simple toggle. When triggered, it executes three actions atomically:

01 — Server deactivated

The MCP server is immediately taken offline across the entire cluster.

02 — All tokens revoked

Every connection token is invalidated. Total lockout — reconnection blocked until new tokens are issued.

03 — WebSocket connections killed

Active connections terminated via Redis pubsub broadcast. Propagates to every runtime node in the cluster.

Full Visibility. Zero Guesswork.

The Vinkius cloud dashboard includes a full MCP Governance suite — real-time analytics and security controls for production AI operations.

Control Plane

KPI dashboard with request volume, latency, success rate, token consumption, and AI-generated operational briefings.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

Incident Log

Error trends, failure rates, status-code breakdowns, and forensic audit trail access.

Get started at cloud.vinkius.com — connect your AI agent in under 60 seconds.

E2B MCP

3 tools available

Cloud-hosted on Vinkius

E2B is the leading sandbox platform for AI code execution, and this MCP connects your agent to it. Instead of letting an AI run commands on your desktop, you can spin up a dedicated Linux environment that lives in the cloud. It handles the heavy lifting of isolation, meaning your agent can test scripts, run data analysis, or execute shell commands in a sandbox that gets destroyed once the job is done. You'll get environments that start almost instantly, whether you need a basic setup or something pre-configured with Python or Node.js. It's a standard part of the modern AI stack for anyone who needs to move beyond just generating text to actually performing tasks. You can find this and other high-quality connectors in the Vinkius catalog to build out your agent's capabilities quickly. By using this MCP, you remove the need to manually set up environments or worry about security vulnerabilities. Your agent gets a clean slate every time it needs to run a script, and you get peace of mind knowing that the execution is completely isolated. Whether you're building a tool for students or a pipeline for data scientists, E2B provides the infrastructure to make your agent functional in the real world.

Core Capabilities

01 — Spin up isolated Linux environments

E2B creates a private Linux VM for your agent to use as a workspace.

03 — Permanently delete finished environments

E2B wipes out sandboxes and their files to keep your costs down.

05 — Execute shell commands securely in the cloud

E2B provides a safe way for your agent to interact with a terminal.

02 — Monitor active sandbox statuses

E2B lets you see every environment currently running in your account.

04 — Run Python scripts in pre-configured environments

E2B handles the setup for Python so your agent can run data scripts immediately.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/e2b — connect your AI agent in three steps.

- 01 Connect your E2B API key to the MCP.
- 02 Tell your agent to create a sandbox with a specific template like python3.
- 03 Your agent runs the code and returns the output to you.

The bottom line is E2B gives your agent a secure, disposable workspace to execute any code it generates.

Built For

E2B is for the developer who is tired of letting an AI run scripts on their local machine.

AI Agent Developer

Testing generated code safely without exposing the host system to arbitrary commands.

Data Scientist

Running analysis scripts in isolated, reproducible cloud environments for consistent results.

EdTech Engineer

Providing students with a safe way to execute code in a web browser without infrastructure risks.

What Changes When You Connect

- 01 Move code execution out of your local environment and into isolated Firecracker microVMs using create_sandbox.
- 02 Get nearly instant results with 150ms cold starts when your agent needs to jump into a new task.
- 03 Keep your infrastructure safe because each sandbox has its own kernel and filesystem, completely separated from your data.

-
- 04 Save on costs by using `kill_sandbox` to wipe out environments the moment your agent finishes its task.
-
- 05 Ensure reproducibility for data science projects by using pre-configured templates like `node` or `python3`.
-

Real-World Applications

Safety first for local scripts

An agent writes a script to clean up files. Instead of running it locally, it uses `create_sandbox` to run it in a sandbox first.

Secure student labs

A platform lets students write code. The agent executes it in a sandbox so students can't mess with the host system.

Automated data analysis

A user wants to process a CSV. The agent spins up a `python3` environment, runs the analysis, and reports the summary.

Isolated web app testing

An agent needs to test a web app. It creates a sandbox, installs dependencies, and runs the test suite.

Patterns to Avoid

Running code locally

X AVOID

Trying to let an agent run shell commands on your laptop.

✓ INSTEAD

Use `create_sandbox` to move the execution to a cloud VM.

Leaving sandboxes open

X AVOID

Forgetting to shut down environments after a task is done.

✓ INSTEAD

Use `kill_sandbox` automatically after the agent finishes its task.

Using the wrong template

X AVOID

Creating a base sandbox for a python task.

✓ INSTEAD

Specify the python3 template in create_sandbox to save setup time.

The Right Fit

Use this if you need a way to let your AI agent actually execute code, run scripts, or interact with a terminal without compromising your security. It's perfect for building autonomous agents that do work rather than just talk. Don't use it if you just need to generate text or if you're looking for a persistent database; this is a disposable, isolated environment for execution. If you need a long-running server with persistent storage, look for a dedicated cloud compute provider instead of a sandbox.

E2B Secure Sandboxing for AI Code Execution

Right now, if you want an AI agent to actually do something with code, you're usually stuck in a loop of copy-pasting. The agent writes a script, you copy it, you paste it into your terminal, and you hope it doesn't break anything or delete your home directory. It's a slow, manual process that feels like a constant security gamble.

This MCP changes that by giving your agent its own hands. Your agent can now spin up a Linux environment, run the script, and give you the results in one go. You stop being the middleman for copy-pasting and start being the person who just checks the final output.

E2B Cloud Sandboxes for Isolated Data Science

Data science teams often struggle with the it works on my machine problem. Setting up the right versions of Python, Node, and system libraries across different machines is a headache that eats up hours of engineering time every week.

With this MCP, you can ensure every run happens in a clean, identical environment. You get guaranteed reproducibility for your scripts without the manual setup, making your agent a much more reliable tool for heavy lifting.

E2B 3 Tools for Secure Code Execution

Use these tools to create, monitor, and delete cloud sandboxes for your AI agent.

#	TOOL	DESCRIPTION
01	kill_sandbox	Use kill_sandbox to permanently delete a sandbox and its files. This releases resources and helps you keep your costs down.
02	create_sandbox	Use create_sandbox to spin up a new isolated cloud sandbox. This provides your agent with a fresh Linux environment that starts in 150ms.
03	list_sandboxes	Use list_sandboxes to see all active environments in your account. It helps you monitor resource usage and keep track of what's running.

See It in Action

Real prompts you can use once this MCP is connected to your AI agent through Vinkius Cloud.

U Create a Python sandbox so I can run a data analysis script.



Sandbox created successfully!

SANDBOX ID	TEMPLATE	STATUS	INFO
sbx_abc123	python3	running	Timeout: 300s

Your sandbox is ready. It's an isolated Linux environment with Python 3.11 pre-installed. Use the sandbox ID to execute code.

U Show me all my running sandboxes.



Found 3 active sandboxes:

SANDBOX ID	TEMPLATE	STATUS	INFO
sbx_abc123	python3	running	Started: 5min ago
sbx_def456	node	running	Started: 12min ago
sbx_ghi789	base	running	Started: 1h ago

U Kill sandbox sbx_ghi789 — I'm done with it.



Sandbox terminated successfully!

SANDBOX ID	STATUS	INFO
sbx_ghi789	killed	Sandbox terminated — resources released

Frequently Asked Questions

01 What does the E2B MCP actually do?

E2B provides a secure way for your AI agent to execute code in the cloud. It creates isolated Linux environments so your agent can run scripts without touching your personal computer.

02 Can I use the E2B MCP to run Python scripts?

Yes, you can use pre-configured templates to spin up environments that already have Python installed. This makes it easy for your agent to run data analysis or scripts immediately.

03 Is the E2B MCP safe for running untrusted code?

It is designed specifically for security. Each sandbox is an isolated microVM with its own kernel and filesystem, meaning the code can't escape to your main system.

04 How fast does E2B MCP spin up a sandbox?

The sandboxes are built for speed and typically start in about 150ms. This allows your AI agent to jump into a new task almost instantly.

05 Can I use the E2B MCP for Node.js projects?

Yes, there is a dedicated Node.js template. Your agent can spin up a workspace with Node pre-installed to run JavaScript applications or test your code.

06 How do I manage my costs with the E2B MCP?

You can manage costs by using the kill tool to delete sandboxes as soon as your agent finishes its task. This ensures you only pay for the time your code is actually running.

07 How secure are E2B sandboxes?

E2B sandboxes run as dedicated Firecracker microVMs — the same technology used by AWS Lambda and Fargate. Each sandbox has its own Linux kernel, filesystem, and network stack, providing hardware-level isolation. Code running in a sandbox cannot access your host system, other sandboxes, or any external resources unless explicitly configured.

08 What programming languages are supported?

E2B supports Python, JavaScript/TypeScript, R, Java, and Bash out of the box. You can also create custom sandbox templates with any pre-installed tools, libraries, or system dependencies. The base template provides a full Ubuntu Linux environment where you can install anything via apt or pip.

09 How does E2B pricing work?

E2B uses usage-based pricing billed per second of compute time. The free Hobby plan includes a one-time \$100 credit (no credit card required), up to 20 concurrent sandboxes, and 1-hour maximum session length. The Pro plan starts at \$150/month with 24-hour sessions and higher concurrency limits.

Go Live in 60 Seconds

Get your connection token from cloud.vinkius.com, then paste the endpoint URL into any MCP-compatible client.

YOUR MCP ENDPOINT

`https://edge.vinkius.com/[TOKEN]/mcp`

CLIENT	WHERE TO CONFIGURE
 Claude AI	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 Cursor	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 VS Code	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"e2b": { "url": "..." }</code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 ChatGPT	Settings → Tools & plugins → Add MCP server → Paste endpoint
 Gemini	Extensions → Add MCP Server → Paste endpoint URL

ASK AN AI ABOUT THIS

Let your preferred AI explain this MCP server

-  Ask ChatGPT 
-  Ask Claude 
-  Ask Perplexity 
-  Ask Gemini 
-  Ask Grok 

READY TO CONNECT

E2B is live on Vinkius Cloud.

Get your connection token, paste it into your AI agent, and start building. No SDK. No deployment. Just results.

Start at cloud.vinkius.com →

vinkius.com · support@vinkius.com

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by E2B. All third-party trademarks, logos, and brand names are the property of their respective owners. Their use in this document is strictly for informational purposes to identify service compatibility and interoperability.

DOCUMENT INFORMATION

Generated	July 2026
MCP Server	E2B MCP
Server ID	019d758b-1ee8-73d2-bcb9-fc66121d9410
Platform	Vinkius Cloud for AI Agents
Endpoint	<code>https://edge.vinkius.com/{token}/mcp</code>

LICENSE & USAGE

This document is generated automatically by the Vinkius PDF Engine. Content reflects the MCP server configuration at the time of generation and may change as updates are deployed. For the most current information, visit vinkius.com/mcp/e2b.