

MCP SERVER

NO CODE

CLOUD HOSTED

ElevenLabs MCP

Generate hyper-realistic speech from any text input.

ElevenLabs MCP connects advanced neural audio synthesis and voice cloning directly into your AI workflow. Convert any text into hyper-realistic speech using dozens of cloned or standard voices, manage your entire audio library, and track usage—all without leaving your agent environment.

A+ Quality Score 98.33/100

elevenlabs

text-to-speech

voice-cloning

tts-api

audio-generation

neural-audio



The infrastructure that powers AI agents in the real world.



Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

Your AI Connections Run Through Vinkius Cloud

The world's largest
managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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.

ElevenLabs MCP

12 tools available

Cloud-hosted on Vinkius

This connector gives you complete control over high-fidelity audio creation from within your AI client. You stop relying on web interfaces and manual copy-pasting for voiceovers. Instead, your agent uses the ElevenLabs tools to orchestrate speech generation: convert text using selected models or clone custom voices against a specific script. Need to check how many characters you've used this month? Your agent can retrieve that data instantly. You also get visibility into all past creations through history tracking and even download direct links for your finished audio files. Because Vinkius hosts the entire catalog, you keep everything—from monitoring account usage to deleting old recordings—in one spot with any MCP-compatible client.

Core Capabilities

01 — Generate speech from text

Send plain text and select a voice or model, and the MCP converts it into high-quality audio.

02 — Access and manage voices

Retrieve details about all available voices to know which one to use for a project.

03 — Check account usage limits

Query your subscription status to see remaining characters or credit balance.

04 — View and manage history

List past audio generations and retrieve the download link for a specific recording.

05 — Fine-tune voice settings

Get detailed information about existing voices to ensure they match your branding guidelines.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/elevenlabs-alternative — connect your AI agent in three steps.

- 01 First, your agent calls the MCP to list all available audio models or current voices.
- 02 Next, you provide the specific text and select a voice ID (or model) for conversion.
- 03 The MCP executes the synthesis and returns metadata, which includes the unique download link for the generated audio.

The bottom line is that your AI client acts as the dedicated interface, handling all API calls and data retrieval so you don't have to touch the web portal.

Built For

This MCP helps content creators and marketing teams who are tired of switching between their editing software, a separate account dashboard, and the cloud editor just to make an audio file. It's for anyone whose job requires massive amounts of high-quality voiceover work.

Video Content Creator

Needs to quickly generate multiple versions of video scripts using different cloned voices and track which ones worked best.

Marketing Automation Specialist

Automates the creation of personalized audio messages for campaigns, while simultaneously tracking character usage against a monthly budget.

Developer/Engineer

Integrates text-to-speech and voice cloning into an application's backend without writing complex external API wrapper code.

What Changes When You Connect

- 01 You eliminate web portal friction. Instead of copy-pasting into a browser editor, your agent runs the `text_to_speech` tool directly, treating audio generation like any other function call.
- 02 Voice management becomes programmatic. Use `list_voices` and `get_voice` to programmatically check which voices are available before making an API call, ensuring you use the right tone for the job.
- 03 Tracking usage is instant. The ability to run `get_subscription_info` means your automation can halt or warn you when you hit credit limits, preventing unexpected failures mid-campaign.
- 04 History management is streamlined. You don't have to manually search; running `list_audio_history` gives you a clean log of all past jobs, and `get_download_link` pulls the file instantly.
- 05 Complex workflows are simplified. Your agent can select the perfect model by calling `list_models`, ensuring your content uses stable and appropriate neural audio synthesis for multilingual needs.

Real-World Applications

Updating a podcast series with new hosts

The producer asks their agent to generate the intro script using the 'Narrator' cloned voice, then asks the agent to list all other available voices so they can select the next host for the following week's episode. This sequence of actions handles both content generation and asset discovery.

Building a dynamic onboarding guide

A developer connects their agent and prompts it to read all system documentation (text) into audio format, using a specific corporate voice profile. The agent then uses `get_voice` to confirm the correct Voice ID is used before running the final `text_to_speech` command.

Auditing marketing spending

The campaign manager asks their agent, 'How much credit did we spend this month?' The agent runs ``get_subscription_info``, giving an immediate answer on remaining characters. Then they use ``list_audio_history`` to see exactly which campaigns consumed the most credits.

Cleaning up old voice assets

A user realizes a test clone is no longer needed. They ask their agent to list all cloned voices using ``list_voices``, identify the unused ID, and then use ``delete_voice`` to remove it entirely from the system.

Patterns to Avoid

Treating it like a simple API wrapper

✗ AVOID

Just calling 'generate audio' without telling your agent which voice, model, or history details you need. This leads to vague errors and requires manual follow-up.

✓ INSTEAD

Always structure the request by first using ``list_voices`` or ``list_models`` to gather necessary IDs, then feeding those specific parameters into the ``text_to_speech`` tool for reliable execution.

Ignoring usage limits

✗ AVOID

Running a large batch of audio generation jobs and getting a cryptic 'quota exceeded' error mid-process because no one checked the balance.

✓ INSTEAD

Start every major workflow by calling ``get_subscription_info``. This gives you real-time data on your remaining characters so you can adjust the scope before running ``text_to_speech``.

Not managing assets

✗ AVOID

Leaving dozens of unused or old cloned voices in the account, making it hard to find the right one when a new project starts.

✓ INSTEAD

After confirming you don't need an old asset, use ``list_voices`` to confirm its ID, and then run ``delete_voice``. It keeps your library clean and functional.

The Right Fit

Use this MCP if the core of your workflow requires repeatable, programmatic audio generation from text. You need the ability for your agent to not just *generate* sound, but also to manage the entire lifecycle: check quotas (`get_subscription_info`), select assets (`list_voices` , `list_models`), track usage (`list_audio_history`),

and retrieve final files (`get_download_link`). Don't use this if you only need simple text-to-speech generation once a month. For that, a basic cloud service might suffice. However, if your process involves advanced voice cloning or multilingual support, this MCP is necessary because it exposes the deep control needed to manage those specific assets and parameters.

Manual audio work is a nightmare of switching tabs.

Today, making an audio file means jumping between your video editor, logging into the ElevenLabs web dashboard, copying text blocks, selecting a voice from dropdowns, hitting 'synthesize,' and then waiting for the download link to appear. If you change voices or need to check usage mid-process, it's another click, another tab, and more copy-pasting.

With this MCP, your agent handles all of that behind the scenes. You just tell your client what text needs conversion and which voice ID to use. Your AI acts as a dedicated producer, running the necessary tools—like `text_to_speech` and `list_voices`—and giving you the final audio link without you ever leaving your workspace.

Get granular control over every asset with ElevenLabs MCP.

Manual systems make it impossible to audit everything. You can't easily check how many characters were used across five different projects, nor can you programmatically delete a single test voice when the project ends. Everything is siloed in web portals or complicated dashboards.

Now, your agent gives you full visibility. By combining `get_subscription_info` with `list_audio_history`, you get an immediate, auditable record of everything that happened and exactly how much budget was consumed. It's total control over your creative output.

ElevenLabs With 12 Tools

Use these tools to scale every part of your audio pipeline, from listing available models to generating final download links for completed speech.

| # | TOOL | DESCRIPTION |
|----|------------------------------------|--|
| 01 | <code>delete_voice</code> | Deletes an uploaded or cloned voice profile to keep your library clean. |
| 02 | <code>get_download_link</code> | Generates a direct, usable URL for any completed audio recording. |
| 03 | <code>get_account_info</code> | Pulls general user profile details associated with your ElevenLabs account. |
| 04 | <code>get_history_item</code> | Retrieves all details about a specific past audio generation job. |
| 05 | <code>get_subscription_info</code> | Checks your current credit usage and displays your plan's limits for the month. |
| 06 | <code>get_voice_settings</code> | Retrieves advanced parameters related to how a specific voice was fine-tuned. |
| 07 | <code>get_voice</code> | Gets detailed metadata for any specified voice profile, including ID and type. |
| 08 | <code>list_audio_history</code> | Returns a list of all past text-to-speech jobs you've run. |
| 09 | <code>list_models</code> | Lists all available neural audio models (like Multilingual v2) for selection. |
| 10 | <code>list_voices</code> | Retrieves a full inventory of every voice you have access to, standard or custom cloned. |
| 11 | <code>text_to_speech</code> | Converts provided text into high-fidelity audio using the selected model and voice. |
| 12 | <code>delete_history_item</code> | Removes a specific record from your audio generation history. |

See It in Action

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

U Convert this text to speech using 'Bella' voice: 'Welcome to the world of neural audio'.



Audio generation triggered! I've converted your text using the 'Bella' voice (ID: EXAV...vXU). You can find the record in your history or I can retrieve the download link for you.

U List all my available voices in ElevenLabs.



I've retrieved your voice library. You have access to 10 standard voices and 2 custom cloned voices ('CEO Voice' and 'Narrator'). Which one would you like to use?

U Check my remaining character limits for this month.



Scanning subscription... You have 15,420 characters remaining out of your 30,000 monthly limit. Your next reset is in 12 days. Would you like to view your generation history?

Frequently Asked Questions

01 How do I check my credit balance using the ElevenLabs MCP?

Run ``get_subscription_info``. This tool immediately pulls your current usage and remaining character count for the month, so you never run out of budget unexpectedly.

02 Can the ElevenLabs MCP clone a voice from my own recordings?

Yes. The MCP allows you to manage cloned voices through tools like ``list_voices`` and ``get_voice``, enabling you to maintain high-fidelity branding across all content.

03 What if I want to delete an old audio recording from my history?

You can use the `delete_history_item` tool. This allows you to programmatically clean up specific records, keeping your workflow organized and focused on current work.

04 Which text-to-speech tool should I use in the ElevenLabs MCP?

Always use `text_to_speech`. This is the primary function that takes input text and synthesizes the final audio output using your chosen voice and model.

05 Does the ElevenLabs MCP support multilingual content?

Yes. By listing models, you can select advanced options like Multilingual v2, ensuring accurate speech synthesis across dozens of languages.

06 How do I find my ElevenLabs API Key?

Log in to your account, click your profile icon (bottom left), and navigate to the API Key section to generate or copy your token.

07 Which model should I use for multiple languages?

The `eleven_multilingual_v2` model is recommended for high-quality speech generation in over 29 different languages.

08 Can I get a direct download link for a past generation?







Yes! Use the `get_download_link` tool with a history item ID to retrieve a temporary URL for the audio file.

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>"elevenlabs-alternative": { "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

ElevenLabs 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](https://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 ElevenLabs. 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 | June 2026 |
| MCP Server | ElevenLabs MCP |
| Server ID | 019dd0e8-8f9a-7192-8383-8a1eeb2ccbe0 |
| Platform | Vinkius Cloud for AI Agents |
| Endpoint | https://edge.vinkius.com/{token}/mcp |

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/elevenlabs-alternative.