MCP SERVER | NO CODE | CLOUD HOSTED

ElevenLabs Alternative MCP

Generate studio-quality audio from any conversation.

ElevenLabs Alternative MCP lets you generate high-fidelity speech and audio effects directly through your AI client. Convert text into lifelike voiceovers, clone existing voices for character consistency, or clean up noisy recordings—all from a single conversation. It manages entire dubbing projects across multiple languages and even generates unique sound effects just by describing them.

A+ Quality Score 98.33/100

text-to-speech | voice-cloning | generative-audio | speech-synthesis | ai-voice | audio-processing

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/elevenlabs-alternative-1

SCAN TO VIEW

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.

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.

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.

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	Ed25519	24h TTL	V8 Isolate
Encryption at rest	PKI vault signatures	Ephemeral session keys	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.

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.

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.

06 — Honeytoken Trap System

Phantom credentials are injected into isolated environments. If a honeytoken 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.

Agent Activity

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

FinOps

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

Tool Health

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

Firewall & DLP

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

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

34 tools available
Cloud-hosted on Vinkius

You don't need to switch between ten different tools to create professional audio content anymore. This MCP lets you treat your AI client like an audio studio. Need a voiceover for a script? You can convert text into speech using custom or pre-built voices, giving granular control over the resulting tone and style. Want to make a character sound consistent across a whole series? Use our voice design tools to create and save unique digital identities. It goes deeper than simple narration: you can transform audio from one voice style to another while keeping the emotion intact, or isolate recordings to strip out background noise completely. Plus, it handles large projects, managing dubbing workflows for multiple languages. By connecting this MCP through Vinkius, your agent gains access to professional-grade audio synthesis and sound design tools without you ever leaving your preferred workflow.

Core Capabilities

01 — Generate Text to Speech

Converts any written text into natural, high-quality speech using a selection of custom or pre-set voices.

02 — Clone and Manage Voices

Designs new unique voice profiles from scratch or finds similar existing voices within your library.

03 — Change Voice in Existing Audio

Transforms an audio recording's vocal style from one character to another while maintaining the original emotion and delivery.

04 — Clean Up Noisy Recordings

Processes existing audio files to remove background noise, resulting in clean voice tracks ready for publishing.

05 — Create Sound Effects on Demand

Generates unique sound effects—like laser blasts or footsteps—simply from a descriptive text prompt.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/elevenlabs-alternative-1** — connect your AI agent in three steps.

01 Subscribe to this MCP and enter your ElevenLabs API Key.

02 Directly ask your AI client to perform an audio task, like generating speech or cleaning up a recording.

03 The agent executes the necessary function using the toolset, providing you with the finished, high-quality audio file.

The bottom line is that you get studio-grade audio production capabilities built right into your chat interface.

Built For

This MCP is for content creators and developers who are tired of juggling separate voice synthesis tools, video editors who need consistent character voices across multiple platforms, or sound designers who waste time cleaning up raw audio tracks. If your job involves *any* kind of high-quality spoken word or background sound, this is for you.

Video Editor

Needs to generate voiceovers for explainer videos quickly and consistently, often requiring multiple languages or character voices.

Podcast Producer

Spends time cleaning up interview audio tracks or creating specific sound effects (e.g., transitions, atmospheric sounds) that need to match the show's tone.

Game Developer

Requires prototyping unique sound effects and character dialogue quickly for testing purposes, saving time on specialized audio asset creation.

What Changes When You Connect

01 You don't have to switch between multiple platforms. By connecting this MCP, your AI client handles everything—from converting text into speech using `create_speech` to managing complex dubbing projects with one command.

02 Character consistency is finally possible. Instead of recording voices manually, you can use `design_voice` or `get_voice` to build and save unique vocal profiles that remain consistent across long-form content.

03 Audio cleanup used to require dedicated software. Now, just ask your agent to run `isolate_audio`, and it strips out background noise from any recording so you get a perfectly clean voice track every time.

04 Sound design gets instant. If you need an explosion sound or footsteps for a game demo, simply prompt the agent to `generate_sound` with a description, eliminating manual asset creation.

05 The system handles advanced transformations too. Use `convert_speech` to instantly change the perceived speaker's voice in an existing recording without losing its original emotion.

Real-World Applications

Creating a multi-lingual training module

A corporate L&D specialist needs to create a product tutorial for five global offices. They ask their agent to `create_dub` the original script into Spanish, French, and German automatically, then use `get_dub_status` to track which language is finished. This saves dozens of hours compared to manual recording sessions.

Podcast cleanup after a bad recording session

A podcast host records an episode with too much room echo and traffic noise. Instead of spending an hour in an audio editor, they tell their agent to run `isolate_audio` on the raw file. The resulting clean track is ready for immediate editing.

Prototyping a video game character

A developer needs quick voice samples for three new enemy types. They prompt their agent to `generate_sound` for 'metallic whirring' and use `design_voice` to prototype three distinct, unique vocal identities before committing to professional recording.

Updating marketing materials quickly

A marketing manager needs a new voiceover for an ad but doesn't have the original talent available. They use `create_speech` with their agent, specifying a pre-existing voice profile and inputting the final script to get instant, high-quality audio.

Patterns to Avoid

Trying to process audio outside of conversation

✗ AVOID

Manually downloading an audio file, opening a separate editor (like Audacity), and running noise reduction filters. This is slow and requires multiple steps.

✓ INSTEAD

Keep the flow in your agent. Instead, ask your agent to run `isolate_audio` directly on the file, letting the system handle the entire cleanup process within the chat interface.

Forgetting voice consistency

✗ AVOID

Generating a series of videos and using different random voices for characters. The final product sounds disjointed and unprofessional.

✓ INSTEAD

First, use `design_voice` to create the character's primary vocal profile. Then, always reference that specific ID when calling `create_speech` so the voice remains consistent across all generated content.

Getting bogged down in API documentation

✗ AVOID

Reading complex manuals and figuring out which combination of parameters (e.g., stability vs. similarity) to use for optimal results.

✓ INSTEAD

Just tell your agent what you need: 'I need a voice that sounds like a calm, deep-voiced historian.' The MCP uses its tools like `find_similar_voices` and `list_voices` behind the scenes to find the best match.

The Right Fit

Use this MCP if your core requirement is high-fidelity audio generation or manipulation. You need to convert text into speech, change voice styles in recordings, or synthesize unique sound effects on demand. If you're building a system that requires consistent character voices across multiple media types (e.g., gaming, education), this toolset gives you the necessary control

points, from `create_designed_voice` to managing full dubbing projects via `list_dubs`. Don't use it if your only need is simple text summarization or data extraction; for that, a standard document processing tool will do. You also don't need this if you just want basic podcast editing and are happy using simple filters—this MCP offers professional-grade isolation and advanced voice cloning capabilities that go far beyond basic cleanup.

The headache of mixing audio tools

Right now, if you need a full video, your workflow looks like this: write the script in one app, record the raw voiceover in a second, clean up the background noise using third-party software, and then send it to an entirely different service just to generate sound effects. You're copying files, switching tabs, and waiting for dozens of disparate services to finish.

With this MCP, you keep all that power inside your agent. You write the script, tell the agent to process it using `create_speech`, and ask it to add background sounds with `generate_sound`—all in one conversation. The audio is built piece by piece right where you're working.

Generating voice identity with ElevenLabs Alternative MCP

The biggest time killer used to be maintaining a character's unique sound across dozens of assets. Every new video meant recording, or at least simulating, the voice from scratch, often leading to noticeable inconsistencies that broke immersion.

Now you can use `design_voice` to build and save that signature tone. The agent manages this profile for you, ensuring that every single speech output references the same core vocal DNA. Consistency is no longer an aspiration; it's a command.

ElevenLabs Alternative with 34 Tools

These tools give your agent granular control over every aspect of professional audio production, from simple text-to-speech to advanced voice cloning.

#	TOOL	DESCRIPTION
01	add_dictionary_from_file	Adds pronunciation rules to the system using a PLS file upload.
02	create_designed_voice	Saves a newly designed voice profile to your library for later use.
03	get_character_stats	Retrieves usage statistics related to character voice models and consumption.
04	create_agent	Sets up and configures a new conversational AI agent for specific tasks.
05	create_dub	Initiates and manages an automated project that translates and dubs content into multiple languages.
06	create_project	Sets up a new, structured studio project to manage related audio assets.
07	create_single_use_token	Generates a temporary token for secure, one-time access to the service.
08	add_dictionary_from_rules	Adds custom pronunciation guides based on specific text rules.
09	isolate_audio	Removes background noise from an uploaded audio file, leaving only the clean voice track.
10	convert_speech	Changes a speaker's voice in an existing audio clip to sound like another character or style.
11	list_projects	Displays a list of your saved studio projects and their current status.
12	list_voices	Lists all the voices available in the library, including custom ones.
13	stream_isolate_audio	Processes audio cleanup and background noise removal in real-time as you stream the file.
14	stream_convert_speech	Performs voice changing (Speech to Speech) on an audio clip while it is streaming through the agent.
15	stream_speech	Converts text into speech in a continuous, real-time stream format for immediate use.

#	TOOL	DESCRIPTION
16	create_speech	Converts plain text input into an audio file using a selected voice and style.
17	delete_history_item	Removes specific items from your usage history log for privacy.
18	delete_voice	Permanently removes a custom voice profile you created or saved.
19	design_voice	Creates an entirely new, unique vocal identity based on a text description or prompt.
20	edit_voice	Modifies parameters and characteristics of an existing custom voice profile.
21	find_similar_voices	Searches your available library to locate voices with tones or qualities similar to a reference audio clip.
22	generate_sound	Creates unique, non-vocal sound effects (like footsteps or explosions) based on text descriptions.
23	get_dub_status	Checks the current progress and status of a multi-language dubbing project.
24	get_history_audio	Downloads the actual audio file for a specific item recorded in your history.
25	get_history_item	Retrieves detailed metadata and information about an item previously generated.
26	get_snapshot	Fetches a saved, read-only version of your current studio project status.
27	get_user	Retrieves basic account information and details about your subscription plan.
28	get_voice	Fetches detailed parameters and metadata for a specific voice ID.
29	list_agent_branches	Lists all available operational branches for your conversational AI agent.
30	list_agents	Retrieves a list of all conversational AI agents you have created or configured.
31	list_dictionaries	Shows all the pronunciation dictionaries currently loaded and available to the system.
32	list_dubs	Displays a list of all dubbing projects you have started or managed.
33	list_history	Provides an overview and list of every audio item your agent has generated in the past.

#	TOOL	DESCRIPTION
34	list_models	Shows all available underlying AI models that can be used for synthesis.

See It in Action

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

U List all available voices in my ElevenLabs library.

I've retrieved your voices. You have 12 voices available, including 'Rachel' (ID: 21m00Tcm4TlvDq8ikWAM) and 'Clyde' (ID: 2EiwWnXFnvU5JabPnv8n). Would you like to use one of these for speech generation?

U Generate a sound effect of a futuristic laser blast.

Generating sound effect... I've created a 'futuristic laser blast' audio. You can now download or play the generated sound effect.

U Convert this text to speech using voice ID pNInz6obpgmqMArWsc7r: 'The future of audio is here.'

Processing text-to-speech... I've generated the audio for your text using the specified voice. The high-quality speech file is ready.

Frequently Asked Questions

01 **How do I generate sound effects with ElevenLabs Alternative MCP?**

You just tell your agent what you want to hear, like 'a cartoon squirrel jumping.' The tool will use `generate_sound` and provide the effect immediately. You don't need to know audio terminology.

02 **Can ElevenLabs Alternative MCP handle multiple languages for dubbing?**

Yes. Your agent manages this with the `create_dub` tool, allowing you to automate translating and generating voiceovers in several different target languages from a single project.

03 **What is the difference between 'designing' and 'creating' a voice?**

Designing uses your prompt to build a brand new unique vocal identity, which you then save with `create_designed_voice`. Creating uses that saved ID when you call `create_speech`.

04 **Does ElevenLabs Alternative MCP let me clean up existing audio?**

Yes. You can use the `isolate_audio` tool to automatically remove background noise, giving you a much cleaner track that's ready for final editing.

05 **How do I keep my characters sounding the same? (ElevenLabs Alternative MCP)**

You must first use `design_voice` to create a unique voice ID. Then, always pass that saved ID into your speech generation calls so the agent maintains character consistency.

06 **How can I convert text to a specific voice using this server?**

You can use the `create_speech` tool. Simply provide the `voice_id` and the `text` you want to synthesize. The agent will generate the audio for you.

07 **Can I see all the voices available in my account?**

Yes! Use the `list_voices` query. It will return a list of all available voices, including their IDs, names, and categories, so you can choose the right one for your project.

08 **Is it possible to remove background noise from an existing audio file?**

Absolutely. Use the `isolate_audio` tool by providing the audio in base64 format. The server will process it and return a clean version with the background noise removed.

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 `"elevenlabs-alternative-1": { "url": "…" }`
Windsurf	MCP Settings → `mcp_settings.json` → 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 →

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	`019e3890-0e3e-715d-9319-ef19c70ef530`
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-1.