

MCP SERVER

NO CODE

CLOUD HOSTED

# ElevenLabs MCP

Produce lifelike audio, from text to multi-lingual video.

ElevenLabs gives your agent full control over AI speech generation, allowing you to create high-fidelity, lifelike voiceovers and dub videos into multiple languages. Manage voices, track usage quotas, and generate complex audio pipelines—all without leaving your chat interface.

**A+** Quality Score 100/100

text-to-speech

voice-synthesis

audio-generation

dubbing

speech-api

natural-language



# 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](https://cloud.vinkius.com) — connect your AI agent in under 60 seconds.

# ElevenLabs MCP

10 tools available

Cloud-hosted on Vinkius

This MCP lets you take the complexity out of professional audio production. Instead of jumping between separate platforms for scripting, recording, and localization, you talk to your agent and it handles the heavy lifting. You can convert raw text into perfect speech using lifelike voices or clone existing ones. If you're working on global content, you don't have to manually manage translation queues; you just tell your agent to dub a video into Spanish or French, and the system initiates the process automatically. Need to know how much budget you have left? You check your usage quota right in the chat. All this power is exposed through Vinkius, giving any MCP-compatible client full access to advanced audio tools that used to require dedicated API coding.

---

## Core Capabilities

### 01 — Generate synthetic speech

Converts written text into high-quality spoken audio using various voices and intonations.

### 03 — Automate multi-language dubbing

Takes existing videos or audio and automatically creates translated versions in dozens of languages.

### 05 — Structure complex speech tasks

Handles detailed audio synthesis jobs by sending specific parameters to the engine for precise control.

### 02 — Manage voice profiles

Browse, retrieve details on, and select from a global library of standard or cloned voices.

### 04 — Monitor usage and billing

Checks your current subscription limits, available character quotas, and past generation history.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/mcp/elevenlabs](https://vinkius.com/mcp/elevenlabs) — connect your AI agent in three steps.

- 01 First, subscribe to this MCP and provide your ElevenLabs API Key to connect your account.
- 02 Next, use a natural conversation prompt to tell your agent exactly what audio you need—for example, 'Dub this video into Italian' or 'Generate speech for X text using voice Y.'
- 03 Finally, the system triggers the necessary pipeline, and you receive confirmation of the job status, tracking ID, or the generated audio file.

The bottom line is that your agent handles all the API calls and complexity; you just talk to it.

---

## Built For

This MCP is for content teams whose workflow hinges on constant, high-quality audio output. You're the marketing manager who needs localized voiceovers across ten countries this week, or the developer building a complex media pipeline that can't afford to wait hours for manual translation.

### Video Content Producer

Needs to take one video and rapidly generate dozens of versions in different languages while maintaining consistent voice quality.

### Developer/DevOps Engineer

Requires a reliable way to test text-to-speech endpoints or validate voice settings directly within an agent workflow without writing boilerplate API code.

### E-learning Course Creator

Needs to update hundreds of training modules with professional, natural-sounding narration and track the associated character usage against a budget.

## What Changes When You Connect

- 01 You get immediate control over voice assets. Instead of guessing which voices work best, you can use the `list_voices` tool to browse and select from a massive library before generating anything.
- 02 Localization becomes trivial. If you need to dub a marketing video into five different languages, your agent handles it with one prompt, managing the entire translation queue for you.
- 03 Quota management is simple. You never run out of budget because you can check your spending limits anytime using `get_subscription`, making content creation predictable.
- 04 Debugging speech becomes easy. If a script has weird pronunciation or missing details, you use `list_pronunciation_dictionaries` to correct it before the audio generation even starts.
- 05 Every job is logged. You can review your work and troubleshoot by calling `list_history`, ensuring nothing gets lost in manual spreadsheets or forgotten folders.

---

## Real-World Applications

### Launching a Global Product Line

A product manager needs to launch the same training video across five countries. Instead of hiring five separate voice actors, they ask their agent to initiate dubbing jobs for all five languages in one go using `list_projects` and `text_to_speech`. This ensures brand consistency and saves weeks of coordination.

### Updating a Technical Manual

A technical writer needs to update speech narration for an existing manual. They first use `list_voices` to select the corporate voice, then run small text segments through `text_to_speech`, and finally compare them against the previous version using `get_history_item`.

### Building a Media Pipeline

A developer needs to build an automated video generation tool. They connect this MCP, allowing their agent to automatically validate voice settings via ``get_voice`` and then use the results to feed into a larger system.

### Auditing Content Spend

A marketing director needs to track which campaigns are eating up the most budget. They ask their agent to call ``list_history``, instantly generating a report on character usage across all departments for easy cost analysis.

---

## Patterns to Avoid

---

### Manual Quota Tracking

#### X AVOID

Waiting until the end of the month when you realize you've exceeded your monthly character limit, forcing an emergency paid upgrade.

#### ✓ INSTEAD

Always check your spending first. Use ``get_subscription`` to see exactly how many characters you have left before starting a large dubbing project.

### Forgetting Voice Details

#### X AVOID

Trying to generate audio with a voice name that was misspelled or deprecated, leading to failed jobs and wasted time.

#### ✓ INSTEAD

Before running ``text_to_speech``, use ``list_voices`` to confirm the exact spelling and availability of the profile you need.

### Losing Job Records

#### X AVOID

Completing a large batch of audio jobs, but having no central place to download or track which file was generated when.

#### ✓ INSTEAD

Immediately use ``list_history`` after a big job. This keeps a clean record and allows you to reference the specific details using ``get_history_item``.

## The Right Fit

Use this MCP if your primary need is generating, managing, or localizing high-quality audio content programmatically. If you are building a system that requires text to become speech—whether for e-learning, media dubbing, or video narration—this is the tool. However, don't use it if you just need simple file conversion (like MP3 to WAV), because this MCP deals with *synthesis*, not format changes. If your problem is connecting audio data to a database record, you might look for a generic data connector instead; this

handles the creative side of speech generation. Remember, calling `text_to_speech` doesn't just generate sound; it also logs metadata and tracks usage, which is key for billing integrity.

---

---

## Manually creating consistent audio across different markets is a nightmare.

Right now, if your company expands internationally, you're stuck in a loop. You get the original script, then you have to copy it into Google Translate or manually send it to human translators. Once translated, you either hire voice actors for each language or spend hours stitching together audio segments from various web services that don't match your brand's core tone.

With this MCP, you give your agent the original content and the target languages. It manages the entire dubbing pipeline automatically, ensuring the correct cross-lingual voices are used and that everything is logged for billing purposes. You get perfect multi-language audio output in a single chat session.

---

## ElevenLabs MCP delivers total control over your synthesized voice assets.

You no longer have to rely on generic, one-size-fits-all voice packs. You can use `list_voices` to identify and select specific cloned or standard voices that match your brand's personality. Furthermore, you don't just generate audio; you also manage the complex parameters using `get_voice`, giving you surgical control over every aspect of the output.

What changes is that the entire voice creation lifecycle—from selection to generation to auditing and quota checking—is now exposed as simple commands in your agent. You get reliability, not just raw audio files.

---

# ElevenLabs: 10 Audio Generation Tools

Use these tools to manage voices, convert text to speech audio, automate global dubbing projects, and keep a detailed record of all your audio generation history.

#	TOOL	DESCRIPTION
01	<code>list_voices</code>	Retrieves a list of all voice profiles available for use in your audio generation projects.
02	<code>list_pronunciation_dictionaries</code>	Shows pre-made dictionaries that help guide the accurate pronunciation of specific words.
03	<code>get_voice</code>	Pulls detailed information about a single voice profile, including its technical parameters.
04	<code>text_to_speech</code>	Converts any block of text into audio and provides associated metadata for the job.
05	<code>list_models</code>	Lists all available AI speech models, letting you choose the best engine for your specific task.
06	<code>list_history</code>	Fetches a list of past audio generation jobs so you can review what was created and when.
07	<code>get_history_item</code>	Retrieves the full details for one specific job from your history, including download links.
08	<code>get_user_info</code>	Gathers general information about your connected ElevenLabs user profile.
09	<code>get_subscription</code>	Checks your current subscription plan and remaining character usage limits.
10	<code>list_projects</code>	Lists all ongoing or completed dubbing and voice projects you've initiated in the system.

---

## See It in Action

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

**U** Generate audio for: 'Hello, this is a lifelike AI voice.' using voice 'abc-123'



Inference pipeline triggered! I've sent your script to ElevenLabs using voice 'abc-123'. The audio is being generated with perfect intonation. Would you like to check the history for the download link?

**U** Show me my remaining character quota



Retrieving subscription info... You have used 15,000 characters out of your 100,000 monthly limit. You have 85,000 characters remaining for this billing cycle.

**U** Dub this video into Spanish: <https://example.com/video.mp4>



Dubbing job initiated! I've started the translation queue for your video into Spanish. I'll provide the tracking ID (dub\_abc) so you can monitor the rendering status.

---

## Frequently Asked Questions

**01** How do I check my remaining character count using ElevenLabs MCP?

You use the `get\_subscription` tool to pull detailed usage information. This instantly shows your current billing cycle status and how many characters you have left for generation.

**02** Can I dub a video into multiple languages at once with ElevenLabs MCP?

Yes, the system handles multi-language queues via the dubbing tools. You simply prompt your agent to translate and synthesize the audio across all necessary target languages.

---

**03 What is the difference between `text_to_speech` and `get_voice` in ElevenLabs MCP?**

`text_to_speech` actually generates the audio from a block of text. In contrast, `get_voice` only pulls metadata and details about a specific voice profile so you know exactly what voices are available.

---

**04 Where can I find my past ElevenLabs jobs?**

You use the `list_history` tool to retrieve an overview of all your previous audio generation activities. From there, you can call `get_history_item` for deep details on a single job.

---

**05 Does ElevenLabs MCP handle complex pronunciations?**

Yes, it does. If specific words are tricky, you use the `list_pronunciation_dictionaries` tool to guide your agent and ensure those words are pronounced correctly in the final audio.

---

# Go Live in 60 Seconds

Get your connection token from [cloud.vinkius.com](https://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": { "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](https://cloud.vinkius.com) →

[vinkius.com](https://vinkius.com) · [support@vinkius.com](mailto: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	019d758f-1705-716f-a10e-f8bef6cb300c
Platform	Vinkius Cloud for AI Agents
Endpoint	<a href="https://edge.vinkius.com/{token}/mcp">https://edge.vinkius.com/{token}/mcp</a>

### 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](https://vinkius.com/mcp/elevenlabs).