

MCP SERVER

NO CODE

CLOUD HOSTED

AudioStack MCP for AI Agents

Automating high-fidelity audio production and speech synthesis

AudioStack lets your AI agents run a complete audio production studio from natural conversation. It generates professional, high-quality speech using over 700 synthetic voices and handles complex mixing and mastering for content creators and ad agencies alike.

D Quality Score 65/100

ai-voice

text-to-speech

audio-production

synthetic-media

audio-mixing



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.

AudioStack MCP

10 tools available

Cloud-hosted on Vinkius

Need to build audio assets? This MCP connects your agent directly to AudioStack, turning simple text commands into finished, polished audio tracks. You can generate studio-quality voiceovers in dozens of languages using a massive library of synthetic voices. It goes way beyond basic text-to-speech; you tell the system what you want—a story, an ad, or a complex soundscape—and it builds the whole thing for you. This capability is crucial for content creators needing rapid asset generation. By connecting AudioStack via Vinkius, your AI agent gains access to professional mixing and mastering tools that handle everything from voice tracks to background music templates. You just talk through the project goals, and the system produces polished audio files ready for distribution.

Core Capabilities

01 — Generate High-Quality Speech

Produce realistic speech recordings using a deep library of over 700 synthetic voices across multiple languages.

03 — Automate Mixing and Mastering

Apply professional industry standards to mixed audio tracks, handling equalization, compression, and final polish automatically.

02 — Compose Complex Audio Productions

Build multi-layered audio files that combine voice, music, and sound effects into one cohesive unit.

One Click on Vinkius — From Prompt to Execution

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

- 01** Subscribe to the AudioStack MCP and enter your unique API Key into your preferred AI client.
- 02** Prompt your agent with a natural language request, such as 'Generate a 60-second ad for coffee using a friendly male voice in Spanish.'
- 03** The agent executes the necessary steps—voice selection, audioform creation, mixing, and mastering—and returns the final audio file URL.

The bottom line is: you skip the manual studio work. You tell your AI what to make, and it handles the entire production chain for professional results.

Built For

This MCP changes the game for anyone who needs to produce audio at scale without hiring a full post-production team. It's essential for ad agencies running localized campaigns, educational content creators building massive course libraries, and developers needing integrated media generation.

Digital Content Creator

Uses the MCP to take scripts and turn them instantly into full audio episodes, complete with background music and mastering.

Ad Agency Producer

Automates localized ad campaigns by generating different voice versions of the same script across multiple languages efficiently.

Application Developer

Integrates dynamic, professional audio generation into an app's backend workflow using simple natural language commands.

What Changes When You Connect

- 01** Scale content output instantly. Instead of spending hours recording voiceovers, you use `text_to_speech` to generate thousands of words across multiple languages in minutes.

- 02 Produce professional broadcast quality assets. The automated mastering function handles mixing and polish that usually requires dedicated studio engineers every time you call `create_mix`.

- 03 Build complex media without code. Use the descriptive structure with `create_audioform` to combine voices, music templates, and effects into a single, cohesive piece.

- 04 Manage your assets efficiently. You can use `list_voices` to quickly find the perfect voice for a script or `list_sound_templates` to pull background music ideas.

- 05 Keep track of everything you make. Use `get_usage_analytics` and `list_media_files` to maintain a clean, auditable record of every piece of content generated.

Real-World Applications

Localizing Global Ad Campaigns

An ad agency needs to run a campaign in five different countries. Instead of booking voice actors, they prompt the agent: 'Generate the same script for all 5 languages using professional male voices.' The system handles multiple calls to `text_to_speech` and ensures consistent tone across every locale.

Developing Interactive Video Content

A video editor needs a trailer that mixes voiceover with specific ambient sounds. They instruct the agent to 'Mix this script (voice) with the forest ambience template.' The system calls `create_audioform` and delivers a polished, single-file asset.

Creating Educational Course Material

A curriculum designer needs a module on particle physics. They ask the agent to 'Create an audio story about quantum entanglement.' The system uses `create_story`, pulling in appropriate sound effects and background music from templates, resulting in a ready-to-use podcast chapter.

Testing Voice Variations for Characters

A game developer needs five distinct character voices. They use `list_voices` to browse options, then run quick tests using `text_to_speech` on a sample line ('Welcome to the village') with each voice until they find the right fit.

Patterns to Avoid

Treating audio generation as simple text output

X AVOID

The user just asks, 'Write me an ad.' and expects a single MP3 file without music or mastering.

✓ INSTEAD

You need to guide the agent's process. First, use ``list_sound_templates`` for background ideas. Then, combine your script with those templates using ``create_audioform`` before running final polish through ``create_mix``.

Ignoring asset management

X AVOID

The user generates dozens of files and quickly loses track of the best voices or mixes.

✓ INSTEAD

Use ``list_voices`` to catalog your options before starting, and always check ``list_media_files`` afterwards. This keeps all your work organized in one spot.

Forgetting the final polish step

X AVOID

The user generates speech and music separately, resulting in a mix that sounds amateur or uneven.

✓ INSTEAD

Don't stop at generation. Always pass your completed tracks through ``create_mix`` to automate professional-grade mixing and mastering.

The Right Fit

Use this MCP if your main bottleneck is the sheer volume or complexity of audio production, not the content itself. If you need to rapidly generate localized ad spots or large libraries of narrated courses, this is your tool. However, don't use it if you require real-time, live recording integration—this handles pre-recorded, studio-quality assets. Similarly, if you only need simple voice transcription from a file, a dedicated audio analysis service is better than using `text_to_speech`. Remember that the MCP excels at composition; for raw data validation or database interactions, look into other specialized AI connectors.

AudioStack MCP for AI Agents: Solving Multi-Language Voiceover Production

Manually creating localized ad campaigns is a huge time sink. You have to copy the script, manually find and book voice actors in each target language (Portuguese, Spanish, French), record them, and then spend hours cleaning up and mastering five separate tracks so they all sound like they belong together.

With this MCP, you simply ask your agent for the task. It handles connecting to the required voices via `list_voices` and generating the entire set of localized audio assets in minutes. You get a folder of perfectly synced, high-fidelity files ready to go.

AudioStack MCP for AI Agents: Mastering Complex Audio Storytelling

Creating narrative content used to require coordinating multiple teams: one for the script, another for sound design, and a third for mixing. The manual steps involved copy-pasting cues into different software programs, ensuring every background track matched the mood of the spoken dialogue.

Now, you write it all out in plain English. You instruct your agent to build the story using `create_audioform`. It pulls pre-mixed templates from `list_sound_templates`, combines them with voices, and delivers a fully mixed, cohesive audio narrative every time.

AudioStack MCP: 10 Tools for Advanced Audio Production

Use these tools to generate speech, mix tracks, create complex storyboards, or manage all your audio assets from one place.

#	TOOL	DESCRIPTION
01	<code>get_voice_details</code>	Retrieves specific details about a chosen synthetic voice for confirmation and usage planning.
02	<code>list_media_files</code>	Shows you all the audio files you've uploaded or generated through your account history.
03	<code>list_sound_templates</code>	Provides an inventory of available music and sound design templates ready for use in a project mix.
04	<code>list_voices</code>	Searches the entire voice library, allowing you to filter by language, gender, or specific provider characteristics.
05	<code>text_to_speech</code>	Converts any given text string into spoken audio using a selected AI voice model.
06	<code>create_audioform</code>	Assembles and generates a fully mixed audio piece by combining multiple elements like music, voices, and sound effects.
07	<code>create_mix</code>	Applies professional mixing and mastering techniques to existing or newly generated audio tracks automatically.
08	<code>create_story</code>	Builds a complete, long-form narrative audio piece optimized for podcasting or serialized content.
09	<code>get_audioform</code>	Checks the status and ultimately retrieves the final URL for an audio production you requested.
10	<code>get_usage_analytics</code>	Provides a metric breakdown of your account's usage history to track costs and consumption.

See It in Action

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

- U** I need an ad for my new coffee blend. Use the 'Friendly' voice in Spanish. Make it 45 seconds long.



Ad Script: El café perfecto.

(Sound of gentle espresso machine running)

Voice: ¿Cansado del mismo sabor? Descubre nuestro Blend Élite, el toque de energía que tu mañana necesita. Desde el primer sorbo hasta el último, te acompaña con notas profundas y un aroma irresistible.

[Mixing applied: Gentle jazz background music fades in and out.]

Status: Complete. Audio file available here. Use `create_mix` if you need to change the genre or tempo.

- U** Can you list all the sound templates related to corporate video intros?



Available Sound Templates: Corporate Intro Pack

TEMPLATE ID	STYLE	DURATION (SEC)	USAGE NOTES
T901	Inspiring Cinematic	20	Best for company mission statements.
T902	Modern Minimalist	15	Use with quick cuts and sharp voice delivery.
T903	Uplifting Orchestral	30	Excellent choice for major product launches.

Please let me know which one you want to use in your next audioform.

U What are the best voices for a serious, documentary-style narration about history?



I filtered the voice library based on depth and professionalism. Here are three top recommendations:

- **Voice ID: v312** (Deep Male): Excellent gravitas; perfect for historical documentaries.
- **Voice ID: v709** (Mature Female): Clear, authoritative tone; ideal for academic lectures.
- **Voice ID: v450** (Neutral/Universal): Highly flexible and consistent; works well regardless of the topic.

Would you like a sample readout for any of these?

Frequently Asked Questions

01 How do I use AudioStack MCP to generate voiceovers in multiple languages?

You simply instruct your agent with the script and the target language. The system handles calling ``text_to_speech`` for each locale, ensuring consistent quality and tone across all versions.

02 Can AudioStack help me mix my own recorded audio files?

Yes. You can upload your tracks and use the MCP's mixing tool to apply professional mastering techniques. It levels out volume, removes background noise, and applies EQ so everything sounds cohesive.

03 Is AudioStack better than just using a basic text-to-speech generator?

Definitely. Basic generators only handle speech. This MCP lets you combine voices with music, effects, and templates into one complex asset—that's where the real power is.

04 What kind of content can I create using AudioStack MCP?

You can build almost anything: educational courses, localized marketing ads, narrative podcasts, or even interactive audio dramas. It's an end-to-end studio in one place.

05 How do I manage all the voices and templates? Is it hard to find what I need?

The MCP provides asset management tools. You can use ``list_voices`` or browse sound templates, making sure you always know exactly which assets are available for your project.

06 If I'm a developer, how do I integrate this into my app?

You connect the MCP to your development workflow. Your agent can then use natural language commands to trigger audio generation and pull the resulting media file directly into your application logic.

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 `"audiostack": { "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

AudioStack 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 AudioStack. 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	AudioStack MCP
Server ID	019d7555-763b-7144-ae46-865b9c201801
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/audiostack.