MCP SERVER | NO CODE | CLOUD HOSTED

Resemble AI MCP

Clone voices, synthesize media, detect deepfakes.

Resemble AI MCP gives you full control over synthetic speech. Generate high-quality audio clips from simple text input, clone voices from recordings, and transform existing speech into any target voice—all through a single connection point. It also includes built-in tools to detect deepfakes and apply digital watermarks, making your media production both powerful and secure.

A+ Quality Score 100/100

voice-cloning | text-to-speech | synthetic-media | speech-synthesis | audio-processing | deepfake-detection

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/resemble-ai

SCAN TO VIEW

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.

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.

Resemble AI MCP

16 tools available
Cloud-hosted on Vinkius

Need to generate professional audio without recording talent? This MCP lets you create and manage synthetic voices directly from your agent. You can turn simple text into high-fidelity audio using custom or system voices, even supporting SSML for fine-tuned control. If you have existing audio, the MCP transforms it, letting you change the voice while keeping the original emotion and timing intact. Keeping track of all your work is easy; you manage projects and keep records organized in one place. Plus, since media authenticity matters, you can detect deepfakes or verify watermarks on any file to guarantee content legitimacy. Connecting this MCP via Vinkius means your AI client—whether it's Claude or Cursor—can handle all these complex audio tasks without needing multiple specialized services.

Core Capabilities

01 — Generate speech from text

You can create new, high-quality audio clips simply by providing text and selecting a voice.

02 — Transform existing voices

The MCP changes an input audio file into a target voice while preserving the original speaker's emotion and rhythm.

03 — Build custom voices

You upload raw recordings to train new, unique voices for your projects.

04 — Manage audio assets

The system lets you organize everything using projects and list all available voice profiles.

05 — Verify media authenticity

You can run checks on an audio file to see if it's synthetic or detect the presence of a digital watermark.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/resemble-ai** — connect your AI agent in three steps.

01 Subscribe to this MCP and provide your Resemble AI API Token.

02 Your agent calls the necessary tools, like creating a new project or listing voices.

03 The platform processes the request—whether it's generating text-to-speech audio or running deepfake detection—and returns the resulting file or data to your client.

The bottom line is that you get advanced voice synthesis and security tools integrated into any workflow, turning complex media tasks into simple conversational commands.

Built For

Content creators who need localized audio for global campaigns. Developers building applications that require automated voice identity transfer. Security professionals needing to verify the source of sensitive audio evidence.

Video Producer

Needs to generate multiple versions of a video's narration quickly, switching voices and languages without hiring talent for every script.

Software Developer

Integrates advanced TTS or deepfake detection directly into an application backend using the MCP tools.

Digital Forensics Analyst

Uses the MCP to analyze suspicious audio files, verifying their source and detecting synthetic manipulation using watermarking checks.

What Changes When You Connect

01 Generate voiceovers instantly. Instead of hiring an actor or recording studio, you use the `create_clip` tool to turn text into professional audio using any available voice.

02 Maintain vocal consistency across projects. Use `speech_to_speech` to transfer a known speaker's unique emotional tone and timing onto new source material, ensuring continuity in your brand messaging.

03 Protect content integrity from the start. You can apply an imperceptible watermark with `add_watermark` and later verify it using `verify_watermark`, proving who created the audio.

04 Stay organized while scaling up. The MCP lets you use `create_project` to group all assets related to one campaign, making it easy to locate everything via `list_projects`.

05 Deepfake defense is built-in. Use `detect_deepfake` on suspicious files to check their source probability, or use the tool when reviewing sensitive media.

Real-World Applications

Localizing a Global Podcast Series

A content team needs to release a podcast in five languages. Instead of coordinating with five different voice actors, they use `create_voice` to clone the host's natural tone and then run `create_clip` repeatedly for each language, keeping perfect vocal consistency across all markets.

Automating E-learning Content

An instructional designer needs hundreds of audio snippets for a new course. They write the scripts, use the MCP's TTS tools to generate every clip via `create_clip`, and then manage all these assets within a dedicated project using `create_project`.

Investigating Media Leaks

A security team receives an anonymous audio file. They immediately use the MCP to run `detect_deepfake`, confirming if it's synthetic, and then run `verify_watermark` to see if any official source protected it.

Updating Character Voices in a Game

A development team needs an NPC character to speak new lines. They use the MCP to clone the original actor's voice using `create_recording`, and then generate the new dialogue using `speech_to_speech` for immediate implementation.

Patterns to Avoid

Treating audio generation as a single step

✗ AVOID

Trying to upload an entire folder of recordings and expecting one command to handle everything. You end up with mixed results because the system needs specific commands for training.

✓ INSTEAD

First, use `list_recordings` to review your source material. Then, you must explicitly call `create_voice` after uploading the necessary data via `create_recording`. Don't skip the voice setup step.

Losing track of assets

✗ AVOID

Generating 50 clips for a campaign and having them scattered across multiple cloud folders with no central index.

✓ INSTEAD

Always start by calling `create_project`. Every new batch of audio, whether generated via `create_clip` or processed using `speech_to_speech`, gets stored and managed within that single project context.

Assuming content is safe

✗ AVOID

Publishing sensitive media without knowing if it was manipulated or who created it, risking brand damage.

✓ INSTEAD

Before publishing any high-stakes audio, run a security check. Use `detect_deepfake` to vet the source and use `add_watermark` on your output for guaranteed provenance.

The Right Fit

Use this MCP if your core need involves generating, manipulating, or verifying synthetic speech. If you are building a system that needs text-to-speech (TTS), voice cloning capabilities, or media provenance checks, this is the right tool. Don't use it if you just need basic audio editing (like trimming silence) — those are general audio processing tools. Also, don't use it if your primary goal is *transcribing* existing speech to text; that requires a dedicated transcription service. This MCP focuses on generation and transformation. If you only want simple file storage, use a standard cloud bucket instead.

The headache of managing voiceovers across multiple platforms

Today, if your brand needs an audio update—say, localizing a message for a new market—you usually have to export the original script and then manually upload it into separate services. You repeat this process for every single language or voice change, wasting hours just managing file names and API calls.

With this MCP, you send one prompt to your agent. It handles the whole chain: creating the project, generating the localized text-to-speech audio using `create_clip`, and organizing it all automatically. You get finished, consistent media ready for deployment.

Generate voice clones with Resemble AI MCP

Before this MCP, changing a speaker's voice meant complex workarounds: recording new audio from the actor, or relying on limited built-in text-to-speech voices that lacked character. It was slow, expensive, and often sounded robotic.

Now, you simply upload training material to generate a custom voice, then use `speech_to_speech` to apply that unique identity to any new script instantly. The quality is high enough that no one can tell it's synthesized.

Resemble AI MCP: 16 Tools for Audio Media

These tools let you manage every aspect of synthetic media, from building custom voices to detecting digital manipulation on uploaded audio files.

#	TOOL	DESCRIPTION
01	add_watermark	Applies an invisible digital watermark signature to protect an audio file's origin.
02	create_clip	Generates a new audio clip from text, supporting advanced SSML formatting.
03	create_project	Sets up a dedicated container for organizing related audio assets and work streams.
04	create_recording	Uploads raw audio files specifically for the purpose of training a new voice model.
05	create_voice	Initiates the process to build and register a brand-new, custom voice profile.
06	delete_voice	Permanently removes a specific custom voice from your available library.
07	detect_deepfake	Analyzes an audio file to calculate the probability of it being AI-generated or synthetic.
08	get_clip	Retrieves the details and content of a specific, previously generated audio clip.
09	get_voice	Fetches comprehensive metadata for a single registered voice profile.
10	list_clips	Provides an overview of all the audio clips stored within a specific project container.
11	list_projects	Retrieves a list of every active and archived project you have set up in the MCP.
12	list_recordings	Shows all the raw audio recordings currently associated with a particular voice profile.
13	list_voices	Returns a comprehensive list of every available custom and system-provided voice for use.
14	speech_to_speech	Transforms an input audio file, changing its speaker's identity to the target voice while keeping the original tone.
15	update_clip	Modifies or revises the content of an existing audio clip within a project.
16	verify_watermark	Checks if a digital watermark is present and valid on an uploaded audio file.

See It in Action

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

U List all my Resemble AI projects and their UUIDs.

I've retrieved your projects. You have 'Marketing 2024' (UUID: proj_123) and 'Game Characters' (UUID: proj_456). Which one would you like to work with?

U Create a new audio clip in project proj_123 saying 'Welcome to the future of voice' using voice voice_789.

Generating clip... Success! The clip has been created in project 'Marketing 2024'. You can access it via UUID clip_abc.

U Analyze this audio URL to see if it's a deepfake: https://example.com/audio.mp3

Running deepfake detection... The analysis is complete. The audio shows a 98% probability of being synthetic (AI-generated).

Frequently Asked Questions

01 **How do I start using Resemble AI MCP for voice cloning?**

You must first subscribe and provide your API token to the MCP. Then, you use `create_voice` and follow up with `create_recording` to upload the necessary source audio.

02 **Can Resemble AI MCP handle multiple projects?**

Yes, absolutely. You can call `list_projects` to see all your work areas, and use `create_project` to segment different campaigns or client accounts.

03 **What is the difference between creating a clip and updating a clip using Resemble AI MCP?**

Use `create_clip` when you are generating audio from scratch, usually with new text. Use `update_clip` if the content of an existing piece needs minor revisions or edits.

04 **How do I check if an audio file is a deepfake using Resemble AI MCP?**

Simply use the `detect_deepfake` tool and provide the URL for the suspicious audio. It will return a probability score indicating how likely it is to be synthetic.

05 **Does Resemble AI MCP support SSML tagging?**

Yes, it supports full SSML (Speech Synthesis Markup Language) within the `create_clip` tool. This allows you fine-grained control over pacing and pronunciation beyond basic text input.

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 `"resemble-ai": { "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

Resemble AI 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 →

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by Resemble AI. 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	Resemble AI MCP
Server ID	`019e38e3-bf66-71c3-a4c3-eddad23ca915`
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/resemble-ai.