

MCP SERVER

NO CODE

CLOUD HOSTED

# Play.ht Voice Cloning MCP

Generate voices and clone speech from samples.

Play.ht MCP lets you generate ultra-realistic speech or clone voices instantly using advanced Text-to-Speech engines directly through your agent. Feed it text, adjust the emotion and speed, or provide a short audio clip to create a digital twin of any voice—all from one place.

**F** Quality Score 3.6/100

text-to-speech

voice-cloning

tts

audio-generation

synthetic-media

speech-synthesis



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

# Play.ht (Voice Cloning) MCP

2 tools available

Cloud-hosted on Vinkius

Need professional audio without recording anything? This MCP lets you generate high-quality speech by converting plain text into lifelike voices. You can fine-tune the output, adjusting not just the words but also the speed, emotional tone, and overall quality to get exactly what you need. The biggest time saver is the instant voice cloning capability; simply provide a short audio sample, and your agent creates a unique voice ID that you can use forever. It's built for scale, letting marketers deploy localized content or developers embed realistic character voices into apps. If you manage multiple AI connections, Vinkius centralizes this powerful audio engine, so you connect once and gain access to professional-grade speech synthesis.

---

## Core Capabilities

### 01 — Generate text-to-speech audio

Convert written scripts into streaming, high-quality audio files.

### 02 — Create instant voice clones

Analyze a short audio sample and generate a unique, usable digital copy of that voice.

### 03 — Control vocal performance metrics

Adjust the generated speech's speed, emotional pitch, and quality before outputting the audio.

# One Click on Vinkius — From Prompt to Execution

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

- 01 Subscribe to this MCP and provide your Play.ht User ID and API Key.
- 02 Use your agent to send text for speech generation or upload an audio file sample for cloning.
- 03 The system processes the request, generating a unique voice ID or streaming the final MP3 audio.

The bottom line is that you skip manual recording and tedious audio post-production by letting your AI client handle the entire process.

---

## Built For

This MCP is essential for content creators who need high volumes of voiceovers, developers building interactive apps, or marketers scaling localized campaigns. It solves the pain point of having to hire voice actors or spend hours recording and editing audio.

### Content Creator

Automates video and podcast production by generating consistent, professional voiceovers for scripts without ever needing a microphone.

### Game Developer

Integrates realistic character dialogue into game environments on the fly, giving NPCs voices that match pre-recorded samples.

### Marketing Specialist

Creates personalized audio messages or localized campaign videos quickly by cloning a brand representative's voice for multiple languages.

---

## What Changes When You Connect

- 01 Never record a voiceover again. You can use the `generate_tts_stream` tool to turn any script into audio instantly, giving you consistent quality for all your video content.

- 02 Build character depth in games or apps using instant cloning. The `create_instant_voice_clone` tool lets you create digital twins of voices from just a short sample clip.
- 03 Maintain brand consistency across global campaigns. Clone a specific voice once and use that unique ID to generate localized audio in multiple languages, saving massive time.
- 04 Fine-tune the emotional impact of your words. Beyond language support, this MCP lets you control emotion, speed, and quality for perfect vocal performance every time.
- 05 Streamline complex workflows. By connecting via Vinkius, you route both the voice models and audio data directly through natural conversation with your agent.

---

## Real-World Applications

### Creating a multilingual educational series

A curriculum designer needs to record a science lesson for English, French, Spanish, and Mandarin. Instead of hiring four voice actors, they use the MCP to generate the same script in all four languages using the cloned voice ID, ensuring tonal consistency across global markets.

### Scaling personalized marketing emails

A sales manager needs to send customized, professional audio greetings to 50 clients. They use the MCP to clone their own voice and run `generate_tts_stream` on individualized text blocks, making every message sound personal and high-touch.

### Developing a character-driven game NPC

A solo developer needs an NPC that sounds like a specific person from concept art. They use `create_instant_voice_clone` with a reference audio clip, giving the AI agent a unique voice ID to program dialogue into their application.

### Building an audiobook prototype

A self-published author wants to test an audiobook concept. They feed the entire manuscript into the MCP, using advanced TTS controls to manage pacing and emotional inflection for a full-length, simulated read-aloud.

---

# Patterns to Avoid

---

## Treating audio as just text

### X AVOID

Writing out a script and simply asking the agent to 'make it sound good.' This ignores emotional context or required pacing.

### ✓ INSTEAD

Use the MCP's controls. Instead of simple generation, specify emotion, speed, and use ``generate_tts_stream`` while defining the exact vocal performance you need.

---

## Forgetting to clone a voice

### X AVOID

Generating random audio that sounds generic or robotic because you didn't establish a consistent source identity.

### ✓ INSTEAD

First, use ``create_instant_voice_clone`` with your brand's sample recording. This gives you the stable Voice ID required for all subsequent generations.

---

## Limiting output to one language

### X AVOID

Writing a global campaign and only getting English audio, forcing manual dubbing in other languages.

### ✓ INSTEAD

The MCP supports multi-language generation. Clone your voice once and use the text input for multiple locales to create localized content at scale.

---

## The Right Fit

Use this MCP if your core need is generating professional, high-fidelity audio without manual recording or studio time. You're building an application that requires character voices (use `create_instant_voice_clone`), or you manage a content pipeline needing massive text volume converted to speech (use `generate_tts_stream`). Don't use this if your goal is generating *original* music, as the tool only handles voice. Also, don't try to generate audio from scratch without source material; always provide either an initial sample for cloning or the full script text.

---

---

## Audio production used to mean expensive studios and slow turnarounds.

Today, getting a voiceover means coordinating writers, paying voice actors, scheduling studio time, and then spending hours in post-production cleaning up breaths or adjusting the emotional tone. For global content, you repeat that entire process for every single language—it's slow, expensive, and requires multiple hands.

With this MCP, you feed your agent the script and click 'generate.' You get polished audio instantly, with the ability to control speed and emotion right in the prompt. The result is immediate, scalable, studio-quality speech that drastically cuts down production time.

---

## Play.ht Voice Cloning gives you total control over your sound identity.

Previously, establishing a brand voice meant recording multiple takes with different talent, which was inconsistent and expensive to manage. You had to treat every new piece of content as a completely separate audio job.

Now, by using `create_instant_voice_clone`, you establish a permanent digital identity for your brand's voice. Every future piece—whether generated via `generate_tts_stream` or used across different languages—shares that same unique vocal signature.

---

# Play.ht (Voice Cloning) MCP with 2 Tools

Use these two tools to convert text into professional audio streams or create custom voice profiles from short audio samples.

#	TOOL	DESCRIPTION
01	<code>create_instant_voice_clone</code>	Takes an uploaded audio file and generates a unique, usable digital clone of that voice.
02	<code>generate_tts_stream</code>	Converts any block of text into streaming, natural-sounding audio using Play.ht's TTS engine.

---

## See It in Action

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

- U** Generate an MP3 of 'Hello, how are you today?' using voice ID 's3://voice-cloning-zero-shot/...' and the Play3.0-mini engine.



I've initiated the TTS generation using the Play3.0-mini engine. The audio stream is being processed for the text 'Hello, how are you today?' with the specified voice.

- U** Create a new voice clone named 'My Assistant' using this base64 audio sample.



I am uploading the audio sample to create your instant voice clone 'My Assistant'. Once finished, you'll receive the unique Voice ID to use in TTS tasks.

- U** Convert this text to speech with a happy emotion and 1.2x speed: 'I am so excited to meet you!'



Generating audio with 'female\_happy' emotion at 1.2x speed. The Play.ht engine is now synthesizing the speech for your request.

---

## Frequently Asked Questions

### 01 How does the Play.ht Voice Cloning MCP handle multiple languages?

It supports generating speech in various global languages like French and Spanish. You provide text in a different language, and the engine processes it while maintaining high quality.

### 02 Can I use the Play.ht Voice Cloning MCP for gaming audio?

Yes. Game developers can upload reference clips using `create\_instant\_voice\_clone` to give Non-Player Characters (NPCs) a consistent, unique voice that matches pre-existing character models.

**03 What is the difference between generating audio and cloning voices?**

Generation uses text input with `generate_tts_stream`. Cloning requires an existing audio sample using `create_instant_voice_clone` to create a new, unique voice ID.

---

**04 Does the Play.ht Voice Cloning MCP require me to record my own voice?**

Yes, for cloning purposes, you must provide a short audio sample. This sample is used by `create_instant_voice_clone` to build the digital twin.

---

**05 Is the audio generated from the Play.ht Voice Cloning MCP high quality?**

The engine supports advanced parameters like emotion and temperature, ensuring the output is lifelike and professional enough for commercial use cases.







---

# 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
 <b>Claude AI</b>	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 <b>Cursor</b>	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 <b>VS Code</b>	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"playht-voice-cloning": {   "url": "..." }</code>
 <b>Windsurf</b>	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 <b>ChatGPT</b>	Settings → Tools & plugins → Add MCP server → Paste endpoint
 <b>Gemini</b>	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

# Play.ht (Voice Cloning) 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 Play.ht (Voice Cloning). 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	Play.ht (Voice Cloning) MCP
Server ID	019e38d7-f6da-70eb-b3ab-300c6d206c63
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/playht-voice-cloning](https://vinkius.com/mcp/playht-voice-cloning).