MCP SERVER | NO CODE | CLOUD HOSTED

Speechmatics MCP

Transcribe Audio and Generate Voiceovers, Automatically.

Speechmatics provides high-accuracy audio processing capabilities right in your agent. Transcribe massive amounts of audio files—whether they're podcasts or meeting recordings—into structured text. You can also convert any written script into natural, human-sounding speech using various voices (like Sarah, Theo, and Megan). It handles everything from batch transcription to job management, giving you full control over your audio pipelines.

A+ Quality Score 100/100

speech-to-text | transcription | text-to-speech | audio-processing | natural-language-processing | voice-synthesis

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/speechmatics

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.

Speechmatics MCP

8 tools available
Cloud-hosted on Vinkius

Dealing with raw audio is a massive headache for any workflow. Before this MCP, turning hours of recorded conversation or video content into usable text required specialized software and tedious manual exports. Now, your agent connects directly to Speechmatics through Vinkius, letting you handle advanced audio processing as part of a natural conversation. You can feed it an audio file—via URL or base64—and quickly start a batch transcription job. Need voiceovers for training videos? Just give it the text and tell it which high-quality voice to use. The system manages all the background work, monitoring your jobs until the transcript is ready for you to pull out in JSON or SRT format.

Core Capabilities

01 — Transcribe audio files

Submit large audio recordings and receive highly accurate written transcripts.

02 — Generate synthetic speech

Turn plain text into high-quality, natural-sounding voice audio using multiple character voices.

03 — Manage transcription jobs

Keep track of every processing task, listing recent activity and checking the status of ongoing jobs.

04 — Retrieve completed transcripts

Pull finished transcriptions in various formats like JSON or plain text for immediate use.

One Click on Vinkius — From Prompt to Execution

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

01 First, your agent initiates a request by submitting the audio file (via URL or base64) to create a new job.

02 Next, you monitor the task status using list_jobs and get_job until the transcription is marked as complete.

03 Finally, you retrieve the finished text or subtitles using get_transcript to integrate it into your workflow.

The bottom line is: you tell your agent what audio needs processing, and it handles the entire lifecycle from submission to retrieval.

Built For

This MCP is built for content teams, data analysts, and developers who treat recorded speech as a core asset. If your job involves turning talks, calls, or podcasts into actionable text, this tool saves you from endless copy-pasting and manual QA.

Podcast Producer

Uses the MCP to transcribe raw audio episodes immediately, then uses generate_tts to create voiceover segments for promotional clips.

Developer

Integrates speech recognition into a prototype app, using tools like create_job and get_transcript to handle backend data processing without managing servers.

Data Analyst

Processes hours of recorded client meetings via batch transcription and then uses the output to populate structured databases for reporting.

What Changes When You Connect

01 Batch processing large files is simple. Use create_job to submit multiple hours of audio at once and handle the entire workload without complex scripting.

02 You get professional voice quality for free. The generate_tts tool lets you turn any script into natural speech using voices like Sarah or Theo, perfect for e-learning modules.

03 Monitoring is built in. You never have to worry if a job failed; list_jobs and get_job let your agent track every single step of the process.

04 Output flexibility means less cleanup time. When you pull results with get_transcript, you can choose JSON, SRT subtitles, or plain text.

05 It's secure and auditable. The create_temp_key tool lets your team manage access credentials without exposing permanent API keys.

Real-World Applications

Indexing internal knowledge bases from calls

A Customer Success Manager has a pile of recorded support calls. They ask their agent to use create_job on all the audio files. The system transcribes everything, and then they pull the clean text using get_transcript, immediately feeding it into an indexed search database.

Creating multilingual training materials

An e-learning developer needs to update voiceovers for a new module. They input the script and tell their agent to use generate_tts with Megan's voice, receiving a ready-to-use audio file instantly.

Automating video subtitling

A content creator finishes recording a podcast episode. Instead of manually transcribing it, they ask their agent to use create_job on the MP3 URL and then pull the output using get_transcript in SRT format for immediate upload.

Auditing usage costs

A team lead wants to know how much audio processing has occurred this month. They ask their agent to run get_usage, getting an instant report on account consumption without having to check a separate dashboard.

Patterns to Avoid

Trying to use simple HTTP requests for jobs

✗ AVOID

Manually making multiple calls just to check if the job finished, leading to complex error handling and brittle code that fails when status codes change.

✓ INSTEAD

Let your agent manage the state. Use list_jobs to see recent activity, then rely on get_job to poll for status updates until it's complete before calling get_transcript.

Assuming raw audio is always usable text

✗ AVOID

A developer gets a transcript but finds it contains time stamps and formatting junk that requires hours of manual cleaning in Excel.

✓ INSTEAD

Use the dedicated tools. Transcribe first with create_job, then ask your agent to process the output using get_transcript to ensure you only receive clean, usable text formats like JSON.

Using a general purpose AI for voice generation

✗ AVOID

Asking an LLM to 'make me sound like a robot reading this,' which results in low-quality, monotonous audio that doesn't match brand standards.

✓ INSTEAD

Use the generate_tts tool. This provides access to professional voices (like Jack or Theo) and ensures the output is high quality, ready for production use.

The Right Fit

You should use this MCP if your core problem revolves around turning audio into text, or text into speech. Specifically, if you need reliable batch transcription of large files (use create_job), or if your workflow requires generating professional voiceovers for content creation (use generate_tts). Don't use it if you just need to send a simple message or read data from a structured spreadsheet; those are better handled by dedicated database connectors. If you only need basic file uploading without job management, this MCP is overkill because the tools like list_jobs and get_job provide essential state tracking that simple APIs lack. This is for entire content pipelines.

Sifting Through Hours of Raw Audio Is a Full-Time Job

Today, getting usable text from audio means clicking into a separate transcription service. You upload the file, wait an unknown time for processing, then you log back in to download a messy export. This cycle involves copying data from one system, pasting it into another for cleanup, and hoping all the timestamps didn't break your script.

With this MCP, audio processing happens inside your agent's conversation flow. You just submit the source file, and the system handles the entire wait period in the background. The finished text or subtitles are delivered directly back to you through a simple command.

Get Natural Speech with generate_tts

Writing scripts and then having a voice artist record them is slow, expensive, and requires scheduling. You also have to manually edit the audio file and sync it perfectly to your video timeline.

Now you input your text and tell your agent which professional voice you want using generate_tts. It delivers the finished, ready-to-use audio file in seconds. The friction point—the wait time and manual labor—is gone.

Speechmatics with 8 Tools

These eight tools let you manage every step of advanced audio processing: submitting jobs, tracking status, generating voiceovers, and retrieving clean transcripts.

#	TOOL	DESCRIPTION
01	`create_job`	Starts a new process to transcribe an audio file from a provided source.
02	`create_temp_key`	Generates secure, temporary API keys for client-side access management.
03	`delete_job`	Removes a transcription job from the system if it was started accidentally or is no longer needed.
04	`generate_tts`	Converts specified text into an audio file using high-quality, natural speech voices.
05	`get_job`	Retrieves the current status and specific details for a single transcription job ID.
06	`get_transcript`	Pulls the final, completed text or subtitle file associated with a finished job.
07	`get_usage`	Checks your current billing consumption and usage statistics for the service.
08	`list_jobs`	Shows a list of all recent transcription jobs you have submitted to the system.

See It in Action

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

U Transcribe the audio file at this URL: https://example.com/audio.mp3

I've started a new transcription job (ID: j123abc). I'll use the default English configuration. You can check its status using `get_job`.

U Generate an audio file of Sarah saying 'Welcome to the future of speech technology'.

Generating speech with Sarah's voice... Done! You can now access the synthesized audio for the text provided.

U List my 5 most recent transcription jobs.

Fetching your recent jobs... I found 5 jobs. The most recent one is 'Meeting_Notes.mp3' (ID: j987xyz) which is currently 'completed'.

Frequently Asked Questions

01 **How do I transcribe a large podcast episode with Speechmatics MCP?**

You start by using create_job, providing the audio URL or base64. Your agent monitors its status until it's complete, then you use get_transcript to pull the final text.

02 **Can I generate subtitles with Speechmatics MCP?**

Yes. After a job finishes using create_job, you can retrieve the transcript using get_transcript and specify SRT format for subtitle files.

03 **Is there a way to track my spending on Speechmatics MCP?**

Absolutely. You use the get_usage tool anytime to check your account consumption statistics without leaving your current workflow.

04 **What is the difference between list_jobs and get_job using Speechmatics MCP?**

list_jobs shows a summary of all recent jobs you've run. Use get_job when you know the specific ID of one job and need detailed status updates on it.

05 **Do I need to manage API keys for Speechmatics MCP?**

Yes, but it’s easy. You can use create_temp_key to generate temporary credentials, keeping your main key secure while allowing controlled access for testing or specific integrations.

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 `"speechmatics": { "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

Speechmatics 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 Speechmatics. 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	Speechmatics MCP
Server ID	`019e38f0-fe67-71c5-b58a-73494de646c4`
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/speechmatics.