MCP SERVER | NO CODE | CLOUD HOSTED

Maestra MCP

Automate Media Translation, Voiceovers & Subtitles

Maestra provides automated media workflows, letting your agent handle everything from video transcription and subtitling to global language translation and synthetic AI voiceover generation. Upload a file once via public URL, and instantly get accurate transcripts in 125+ languages, making it ideal for large-scale content distribution.

A+ Quality Score 100/100

transcription | voiceover | speech-to-text | media-processing | automated-translation | synthetic-voice

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

SCAN TO VIEW

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.

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.

Maestra MCP

8 tools available
Cloud-hosted on Vinkius

Your agent connects to Maestra when you need media assets processed at scale. Instead of juggling multiple services or writing complex code just to translate subtitles, this MCP lets your client handle the entire workflow conversationally. You can upload a video file and immediately get speaker-aware transcripts for free. From there, the agent translates those transcriptions into over 125 languages. Need it dubbed? Maestra generates high-quality synthetic voiceovers using multiple AI voices. It's built to manage content distribution—if you need subtitles translated for global reach, this is where your workflow happens. You can also list all your media files or check the processing status on demand. Connecting through Vinkius gives your agent access to this powerful toolset without needing complex setup.

Core Capabilities

01 — Automate Media Transcription

Upload audio or video via a public URL and receive accurate, speaker-aware transcriptions.

02 — Translate Transcripts Globally

Convert existing text transcripts into over 125 different languages using natural language commands.

03 — Generate AI Voiceovers

Create high-quality synthetic voice tracks for any media file using a selection of available voices.

04 — Manage Media Assets

List all content in your account and check the status or details of specific files.

05 — Export Processed Files

Generate temporary download links for results in formats like SRT, VTT, PDF, or JSON.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to the Maestra MCP and provide your API key.

02 Ask your AI client to process a media file using its public URL (e.g., 'Upload this video for transcription').

03 Wait for the initial transcript, then prompt for subsequent steps like translation or voiceover generation.

The bottom line is you talk to your agent in plain language and it handles the complex media pipeline behind the scenes.

Built For

This MCP solves the problem of scaling global content distribution. It's for anyone whose job involves taking one piece of media—a webinar, a documentary, or marketing video—and turning it into dozens of localized assets without manual intervention.

Localization Manager

They use this to automate the entire cycle: getting transcripts from source content, translating those scripts into target languages, and then re-voicing the video for global release.

Content Creator

They rely on it to quickly generate subtitled versions of their videos in multiple languages so they can post evergreen content everywhere without hiring a full translation team.

Developer

They integrate the transcription and dubbing logic into custom applications, using the agent to manage file uploads and status checks programmatically.

What Changes When You Connect

01 Instant Transcripts: Uploading media via `upload_media_for_transcription` gives you accurate, speaker-aware transcripts immediately. You don't wait days for manual captioning.

02 Global Reach: Once transcribed, the agent translates the text into over 125 languages using `translate_transcription`. This lets your content reach massive audiences effortlessly.

03 Professional Dubbing: Need to dub a video? Use `generate_ai_voiceover` with selected voices from `list_available_ai_voices` to create high-quality, synthetic voice tracks for any language.

04 Full Content Lifecycle Management: You can track everything. Use `list_maestra_files` and `get_file_details` to monitor processing status or organize content into folders using `list_account_folders`.

05 Easy Sharing: When done, you get temporary links via `export_transcription_results`, allowing immediate download of subtitles in SRT, VTT, or JSON formats.

Real-World Applications

Launching a Product Globally

The marketing team records an English webinar. Instead of hiring translators for 10 countries, the agent processes the video, generating transcripts and then calling `translate_transcription` repeatedly for all required languages. Finally, it uses `generate_ai_voiceover` to dub the content into local voices.

Archiving Video Content

The development team needs to archive old training videos. They use `list_maestra_files` first to see what's available, then submit new video URLs via `upload_media_for_transcription` so the system can generate searchable transcripts and save them for later reference.

Building a Multilingual App Feature

A developer needs to add support for multiple languages. They use Maestra's tools, specifically `list_available_ai_voices` to pick the right voice type and `generate_ai_voiceover` to ensure all new content has professional, synthetic audio.

Checking Translation Status

A localization team member uploaded a batch of files. They use `get_file_details` to check if the translation job is complete before they try to call `export_transcription_results`. This prevents API errors.

Patterns to Avoid

Treating Maestra like a simple file storage.

✗ AVOID

A user tries to upload raw files directly via the agent without providing a public URL or specifying the source language. The request fails immediately because the tool needs specific parameters to start processing.

✓ INSTEAD

Always provide the required context: Use `upload_media_for_transcription` and ensure you include both the public file URL and the correct target source language.

Trying to manually stitch together different services.

✗ AVOID

A user first transcribes a video using one API, then has to take that text into a second service for translation, and finally copy it into a third system for dubbing. This involves multiple manual handoffs.

✓ INSTEAD

Use the Maestra MCP's structured workflow: Start with `upload_media_for_transcription`, follow up with `translate_transcription`, and finish with `generate_ai_voiceover`—all within one agent conversation.

Forgetting to check file readiness.

✗ AVOID

The agent attempts to generate a voiceover for a media asset that hasn't finished transcription yet. The job fails because the source data isn't ready.

✓ INSTEAD

Before generating, always call `get_file_details` to confirm the processing status is 'complete' and that transcripts are available.

The Right Fit

Use this MCP if your core challenge involves scaling content from one language or format into many others. If you have a media file (video, audio) and need it transcribed, translated, or dubbed for international release, Maestra is the right fit. Don't use it if your goal is simply to index documents for retrieval; use a dedicated

knowledge base tool instead. Similarly, don't rely on it just for simple data storage; always check `list_maestra_files` and `get_file_details` first to confirm the file status. This MCP handles the *content* transformation—the media pipeline itself.

The Problem with Manual Global Content Release

Today, launching a single piece of video content into five different markets means multiple headaches. You have to download the original file, manually upload it to a transcription service, wait for the text, copy that text into a translation tool, and then repeat the whole process—including coordinating separate voiceover recording sessions or finding an expensive dubbing agency. It's slow, error-prone, and costs serious time.

With Maestra, you just point your agent at the media file. The system handles the entire chain: transcription, translation into 125+ languages, and professional synthetic voiceovers—all in one conversation. You get ready-to-use, localized assets without leaving your chat window.

Generating Voiceovers with `generate_ai_voiceover`

Previously, getting a professional voiceover meant hiring talent or using basic, robotic tools. You'd manually select the right tone and gender for every language and have to wait days for the final audio files.

Now, you simply ask your agent to generate an AI voiceover for the processed media. The system handles selecting from all available voices listed by `list_available_ai_voices`, ensuring a high-quality dub that matches your brand's needs instantly.

Maestra: 8 Media Processing Tools

These tools let you manage the full lifecycle of media assets, from uploading and transcribing to translating and generating professional voiceovers.

#	TOOL	DESCRIPTION
01	`export_transcription_results`	This tool generates a temporary download link for all processed file results.
02	`list_maestra_files`	It lists every audio and video file currently stored in your Maestra account.
03	`list_account_folders`	This tool retrieves a list of all content folders set up in your account.
04	`get_file_details`	It fetches the status and specific details for one file by its ID.
05	`translate_transcription`	This function translates an existing transcript into a specified new language.
06	`upload_media_for_transcription`	You use this to submit a new file via public URL and specify the source language for transcription.
07	`generate_ai_voiceover`	It creates a synthetic voice track (dubbing) for an existing media file.
08	`list_available_ai_voices`	This tool shows you all the names and types of AI voices available for selection.

See It in Action

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

U Upload the video at 'https://example.com/video.mp4' for English transcription in Maestra.

I've initiated the upload for your video. Maestra is now processing the file for English transcription. The file ID is 'vid-12345'.

U List all available AI voices for French.

I found 8 synthetic voices for French, including 'Claire (Female)', 'Mathieu (Male)', and 'Amelie (Female)'. Which one would you like to use for your voiceover?

U Get an SRT export link for file ID 'vid-12345'.

I've generated the export link for your subtitles. You can download the SRT file here: https://api.maestra.ai/temp/vid-12345.srt

Frequently Asked Questions

01 **How many languages can Maestra handle?**

Maestra supports transcripts and translations into over 125 different languages. You just need to specify the target language in your prompt.

02 **Does Maestra support video files, or only audio?**

It handles both. The system accepts media files via public URLs for transcription and subsequent processing like translation and voiceovers.

03 **Can I list all my existing videos with Maestra?**

Yes, you can use the `list_maestra_files` tool to get a complete inventory of every file in your account. You can then check specific statuses using `get_file_details`.

04 **What is required to generate an export link with Maestra?**

You must first have completed the media processing (like transcription or translation). After completion, you call `export_transcription_results` to get a temporary download link for your files.

05 **Is there a way to list available voices in Maestra?**

Absolutely. The `list_available_ai_voices` tool lets you see all the synthetic voice options, helping you choose the right tone before running a voiceover job.

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 `"maestra": { "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

Maestra 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 Maestra. 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	Maestra MCP
Server ID	`019d75cb-9a78-71fe-915b-af970aa2a473`
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/maestra.