

MCP SERVER

NO CODE

CLOUD HOSTED

Synthesis MCP

Automate Corporate Video Creation & Dubbing

Synthesis MCP lets you generate professional, corporate AI avatar videos from text prompts directly via your agent. You can programmatically pull templates, list available avatars and voices, or automatically dub existing media into new languages without leaving your terminal.

A+ Quality Score 100/100

ai-video

synthetic-media

avatar-generation

video-dubbing

content-creation



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.

Synthesia MCP

10 tools available

Cloud-hosted on Vinkius

Need to build training modules or internal comms video? This MCP connects your conversational client to the Synthesia platform, turning your AI assistant into a virtual studio director. Instead of logging into a web portal and manually managing renders, you tell your agent what you need. It handles everything from listing available avatars to creating entire videos using pre-set templates with simple data payloads. You can even take an existing video and use the `dub_video` tool to instantly translate it into dozens of languages.

By connecting this MCP through Vinkius, you give your agent operational access to a professional video pipeline. Your workflow goes from concept to final rendered file—all within the chat interface. This means corporate training leads can dispatch bulk instructional videos and localization teams can rapidly scale their video catalogs without ever touching a graphical user interface.

Core Capabilities

01 — Generate Videos from Script

Creates an AI avatar video based on text you provide.

02 — Use Pre-Built Templates

Generates a video by filling in the blanks of a specific, pre-designed template.

03 — Manage Video Statuses

Checks if a video render is finished or deletes old drafts completely.

04 — Find Avatars and Voices

Lists all available professional avatars and voice models for selection.

05 — Automate Localization & Dubbing

Takes an existing video and automatically dubs it into different languages using native voices.

One Click on Vinkius — From Prompt to Execution

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

- 01 Append the Synthesia MCP module to your Vinkius connective logic rules.
- 02 Pass your specific account authorization key (tied to Creator plans or beyond) to the host configuration.
- 03 Tell your agent exactly what to do, like listing all avatars and then generating a video saying 'Welcome' in Portuguese.

The bottom line is you give your AI client the necessary keys; it handles the complex API calls for you.

Built For

Corporate training leads who are tired of managing video renders in a web dashboard. Localization teams drowning in manual translation requests. Internal communications engineers needing to automate high-volume content creation.

L&D Specialist

Needs to programmatically dispatch dozens of multi-lingual training videos using templates and bulk avatar streams.

Global Content Manager

Handles localization, needing to quickly dub existing video content into new languages and manage the source files.

Internal Comms Engineer

Must automate routine internal announcements, generating high-quality avatar videos directly inside an automated pipeline.

What Changes When You Connect

- 01 You eliminate manual platform switching. Instead of jumping between a CMS and the video tool, your agent handles avatar generation and rendering entirely within the chat interface.

-
- 02** Scaling global content is simple. Using `dub_video`, you can take one core training module and instantly generate accurate versions in dozens of languages—no human translator needed for dubbing.
-
- 03** Template consistency matters. With `list_templates` and `get_template_details`, you ensure every video follows corporate branding guidelines, without manual checks.
-
- 04** You maintain control over your assets. You can check the status using `get_video_details` or clear out old drafts instantly with `delete_video` to keep things tidy.
-
- 05** Creating variety is fast. If a script changes, you don't rebuild from scratch; you just use `create_avatar_video` and feed it new text.
-

Real-World Applications

Updating Global Compliance Training

A L&D specialist needs to update a core compliance video for 12 different international offices. Instead of manually re-recording the content twelve times, they run the original video through `dub_video` using the available voice list and get localized masters instantly.

Cleaning Up Old Media Assets

A content manager finds dozens of old draft videos cluttering their account. They simply prompt their agent to list all files via `list_videos`, and then call `delete_video` on the junk, keeping only the finalized work.

Launching New Product Onboarding

An internal comms engineer needs to launch a new product training module. They first use `list_avatars` to pick a professional face, then select a template using `create_video_from_template`, and finally generate the video using `create_avatar_video` all in one go.

Developing Multi-lingual Marketing Assets

A marketing team needs a video promoting a new feature in five languages. They use `list_voices` to select high-quality voices, and then trigger `dub_video` against their master source file for rapid global deployment.

Patterns to Avoid

Treating it like a simple text generator

X AVOID

Trying to just paste a script and expect the video. You have to tell your agent which avatar and template to use, or it fails.

✓ INSTEAD

You must first call `list_avatars` to pick an identity, then decide if you're using raw text (`create_avatar_video`) or filling out a structured format (`create_video_from_template`).

Forgetting the workflow steps

X AVOID

Trying to dub a video without knowing what template it was built on. The system needs context before localization can happen.

✓ INSTEAD

Always confirm your process first by using `get_template_details` before running any generation or modification tools like `dub_video`.

Running multiple simultaneous builds

X AVOID

Sending 50 requests at once without checking status. You won't know which render succeeded and which failed.

✓ INSTEAD

After submitting a video task, always check its progress using `get_video_details` before assuming it's ready to download.

The Right Fit

Use this MCP if your core need is automated, programmatic video asset management. If you routinely generate training materials, internal announcements, or localized marketing content, this is for you. You must be able to define a source script and map that script into structured templates or avatars. Don't use it if you only need simple image generation; this MCP deals exclusively with full video files. Also, don't use it if your workflow requires manual human sign-off on the final rendered media; while we automate production, the final approval step remains yours.

The Pain of Manual Video Production

Right now, creating global training content means copying a master video into a web portal. You manually download that file, upload it to a localization service, wait for the dubbing, and then re-upload it—all across different tabs and platforms. It's slow, tedious, and prone to version control nightmares.

With this MCP, your agent takes over that entire process. Tell it you need an announcement video on 'HR Policy Updates,' specify the desired avatar, and let it handle the full render cycle. You get a ready-to-use file ID without leaving your terminal.

Synthesia MCP: Video Assets On Demand

You no longer have to manually track asset IDs, check status pages, or remember which version you were working on. You simply ask the agent to list all existing videos using `list_videos`, and it gives you a clean manifest.

The difference is speed and consistency. Instead of spending half a day clicking through dashboards, you get instant, programmatic control over your video assets.

Synthesia: 10 Video & Avatar Tools

These tools allow you to programmatically list avatars, create full video renders, automate template filling, and handle international dubbing for all your corporate content.

#	TOOL	DESCRIPTION
01	<code>create_avatar_video</code>	Generates a new video file using a text script and a selected AI avatar.
02	<code>create_video_from_template</code>	Builds a video by populating the variables within an existing, defined template structure.
03	<code>get_video_details</code>	Retrieves the current status and detailed metrics for any specific video render.
04	<code>list_videos</code>	Fetches a list of all videos currently stored in your account.
05	<code>delete_video</code>	Permanently removes a designated video file. This action cannot be undone.
06	<code>list_avatars</code>	Returns an inventory of all professional AI avatars available to use in videos.
07	<code>list_templates</code>	Provides a catalog of video templates you can choose from for creation.
08	<code>get_template_details</code>	Retrieves the specific structure and variable requirements for any chosen template.
09	<code>list_voices</code>	Lists all available AI voice models, including language codes and accents.
10	<code>dub_video</code>	Translates and dubs an existing video into a new target language while keeping the original avatar.

See It in Action

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

- U** Give me a list of internal AI video templates, select one related to 'customer support', and execute a creation pass using proper test inputs.



Accessing Synthesia directories ('list_templates')...
Found template 'Support Starter' (ID: tmpl-1234).
Fetching target blueprint variables natively ('get_template_details').
Dispatching 'create_video_from_template' with payload values successfully...
Video task submitted. The ID is 'vid-5678'. Let me know when you want to execute 'get_video_details'.

- U** List the available AI avatars, focusing on professional corporate styles.



Querying Synthesia avatar registry ('list_avatars')...
Here are some of the active top corporate avatars available for rendering:

- 'Anna_Business' (ID: 101) - Blazer, formal.
- 'Marcus_Exec' (ID: 105) - Suit, authoritative.
- 'Elena_Office' (ID: 112) - Smart casual, friendly.

- U** Check the status of my recent video task 'vid-9920'.



Evaluating task bounds ('get_video_details')...
The render for 'vid-9920' is COMPLETE.
Total duration: 45 seconds.
You can download and stream the final file here: `https://synthesia.io/stream/vid-9920_xyz`.

Frequently Asked Questions

01 How do I generate an avatar video using the Synthesia MCP?

You use `create_avatar_video` and provide a script. This tool generates the video ID which you can then track for completion.

02 Can the Synthesia MCP dub videos into multiple languages?

Yes, using `dub_video`, you can take an existing recorded video and automatically translate and generate versions in different target languages.

03 What if I want to use a specific corporate layout?

Use `list_templates` first to see available layouts, then select one and use `create_video_from_template` with your data payloads.

04 How do I check if my video finished rendering via Synthesia MCP?

Call the `get_video_details` tool. This provides the current status, total duration, and confirmation that the render is complete.

05 Does the Synthesia MCP require me to list avatars first?







No, but it's smart practice. Use `list_avatars` upfront so you know exactly which professional faces are available for your scripts.

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 <code>"synthesia": { "url": "..." }</code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → 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

Synthesia 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 Synthesia. 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	Synthesia MCP
Server ID	019d760f-c6e3-721c-b660-390048474741
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/synthesia.