

MCP SERVER

NO CODE

CLOUD HOSTED

HeyGen MCP

Automate talking head and global video production.

HeyGen automates studio-quality video production, letting you manage AI avatars, templates, and complex videos using simple conversation prompts. Connect this MCP to any agent to list avatars, generate talking head content for social media, or translate existing videos into multiple languages with perfect lip-syncing.

A+ Quality Score 100/100

ai-avatars

video-generation

talking-head

content-creation

studio-quality

template-management



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.

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.

HeyGen MCP

11 tools available
Cloud-hosted on Vinkius

Connect your HeyGen account through this MCP to give your AI client full control over video creation. You can treat your agent like a dedicated Creative Director; it handles the heavy lifting of production management via natural conversation.

Need talking head videos? Your agent lists all available avatars and triggers video generation using specific scenes and avatar configs. Want to scale content for global teams? Use templates by providing variable mappings, generating personalized videos quickly. You can also monitor complex jobs with `get_video_status` to check progress or grab the final URL when it's ready.

If you need to update existing materials, your agent translates a video into multiple languages, retaining natural lip-syncing. Plus, you can list all templates and browse your history of generated videos. By connecting through Vinkius, you get access to this full suite of tools from any compatible AI client.

Core Capabilities

01 — Find available avatars

The MCP lists groups or individual avatars so you know which digital spokesperson to use for your project.

03 — Create videos from templates

The MCP allows you to automate mass content creation by replacing variables within pre-designed video templates using ``generate_from_template``.

05 — Translate videos into new languages

The MCP translates an existing video into another language, automatically syncing the lip movements for realism using ``translate_video``.

02 — Generate video content manually

You can trigger a studio-quality video generation task by providing specific scenes and avatar settings via the ``generate_video`` tool.

04 — Check production status and retrieve links

Use the ``get_video_status`` tool to monitor a job's progress, retrieving the final URL once the video is complete.

One Click on Vinkius — From Prompt to Execution

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

- 01 First, subscribe to this MCP and provide your HeyGen API Key in the agent's settings.
- 02 Next, tell your AI client exactly what you need—for example, 'Generate a video of John Doe using template X.'
- 03 The agent calls the appropriate tool, manages the process, and reports back when the final asset is ready.

The bottom line is that your agent handles the entire production pipeline without you having to manually interact with the HeyGen studio dashboard.

Built For

This MCP is for content creators and marketing managers who are tired of spending hours in a video editor, manually generating videos for every campaign. If your job involves rapid, high-volume communication—whether it's training materials or personalized outreach—you need this.

Content Creator

Generating talking head videos for social media or internal training without ever needing a camera.

Marketing Manager

Automating personalized video outreach to hundreds of clients using pre-built templates at scale.

L&D Specialist

Quickly updating and translating global training modules, ensuring consistent messaging across different languages.

What Changes When You Connect

- 01 Scale personalized outreach. Use `generate_from_template` to automatically populate variables across hundreds of videos, making mass communication effortless.

-
- 02 Never deal with camera setup again. The MCP lets you generate high-quality, talking head content for social media or corporate training using any listed avatar.

 - 03 Manage complex jobs easily. Instead of checking a dashboard, ask your agent to check the status using `get_video_status` and get notified when it's done.

 - 04 Go global instantly. With one prompt, use `translate_video` to turn an English video into Spanish or French while preserving perfect lip movements for local credibility.

 - 05 Understand your options fast. Use `list_avatars` to check every available spokesperson before starting a project, ensuring you pick the right look and tone.
-

Real-World Applications

Rolling out global training content

A learning & development manager needs to update an HR policy video for five different countries. Instead of manually re-recording or coordinating translations, they tell their agent: 'Translate the existing video into German and Portuguese.' The `translate_video` tool handles everything, giving them localized assets immediately.

Rapid social media content

A creator needs daily short tips videos but doesn't have time for the studio interface. They use their agent to first `list_avatars` to pick a specific spokesperson and then call `generate_video` with three scenes, getting the finished asset in minutes.

Personalized sales follow-ups

A marketing team runs a campaign needing 50 unique videos for prospects. They use the agent to call `generate_from_template`, passing in name and company variables, ensuring each prospect gets a highly customized, professional video message.

Tracking campaign assets

A project lead needs to know if last week's high-priority video is ready. Instead of guessing, they ask their agent to check the status using `get_video_status` and get confirmation with a direct download link.

Patterns to Avoid

Manually recreating variables

X AVOID

The user tries to manually pass all variable mappings as raw JSON strings into the agent prompt, which fails because they forget required structure elements.

✓ INSTEAD

Don't write the JSON yourself. Tell your agent: 'Generate a video using template X and replace {'name': 'Sarah'}.' The agent uses ``generate_from_template`` correctly.

Assuming availability

X AVOID

A user tries to generate a video without knowing which avatars are available, resulting in an error because they use an invalid avatar ID.

✓ INSTEAD

Always start by calling ``list_avatars`` first. This confirms the correct `'avatar_id'` before you attempt any generation using ``generate_video``.

Over-reliance on a single tool

X AVOID

The user only uses ``generate_video`` and gets frustrated when they realize the content needs to be localized for Japan.

✓ INSTEAD

Don't stop at generation. After creating the video, you must run ``translate_video`` to adapt it for new markets while keeping the core message intact.

The Right Fit

Use this MCP if your primary need is high-volume, automated video production across different languages and templates. You're here because you want an agent to manage the entire lifecycle: from listing avatars (using `list_avatars`) to creating a draft (`generate_video`), then scaling it via templates (`generate_from_template`), and finally localizing it (`translate_video`). Don't use this if your goal is simple asset storage or file conversion; you need dedicated media management tools for that. Also, don't try to write the complex JSON structures yourself; let the agent handle calling these specific tools so you just talk naturally.

The headache of manual video production workflows

Today, creating a campaign that needs talking heads in three different languages feels like coordinating an army. You jump between the HeyGen studio interface, copy-pasting variable names into templates, manually checking job statuses across multiple tabs, and then—if you want it localized—you have to initiate separate translation jobs for every single piece of content.

It's a clock-watching, click-heavy nightmare. You spend more time managing the workflow than actually creating the message. When you connect this MCP via Vinkius, your AI client steps in and acts as that dedicated Video Producer. It manages all the complex sequencing—from avatar selection to final translation—through simple conversation.

HeyGen provides automated video production with `generate_video``

The biggest time drain is initiating and monitoring single videos. You have to select the avatar, define the scene JSON structure, submit it, and then wait. If you're running 50 variations, that means 50 separate manual steps.

With this MCP, you just tell your agent: 'Generate a video of the new product launch.' It handles calling `generate_video`, structuring the body data, submitting the job, and even reporting back when it's done. The entire process is wrapped into one conversation.

HeyGen: 11 Tools for Video Production

This set of tools lets you manage every aspect of AI video production, from listing avatars to translating final assets.

#	TOOL	DESCRIPTION
01	<code>generate_from_template</code>	Creates a video by replacing variables within an existing template structure.
02	<code>generate_video</code>	Manually starts and triggers a studio-quality video generation task using specific scenes and avatar configurations.
03	<code>get_api_profile</code>	Retrieves basic operational information about the current API account setup.
04	<code>get_template_details</code>	Provides the structure and definitions for a specific video template, letting you know what variables are available.
05	<code>get_video_status</code>	Checks if a generated video is finished and, if so, returns the final download URL.
06	<code>list_avatar_groups</code>	Lists categories of avatars available in HeyGen for selection.
07	<code>list_avatars</code>	Retrieves a full list of all digital avatars, which you need to select an ID for video generation.
08	<code>list_streaming_history</code>	Shows the history of any interactive live streaming sessions conducted through the platform.
09	<code>list_templates</code>	Retrieves a list of all available video templates you can use for content creation.
10	<code>list_videos</code>	Shows a history log of all videos that have been generated via the platform.
11	<code>translate_video</code>	Translates an existing video into another language while maintaining accurate lip-syncing.

See It in Action

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

U List all available avatars and find one named 'Joshua'.



I've retrieved the list of avatars. I found 'Joshua in Suit' (ID: joshua_v3) and 'Joshua Casual'. Would you like to use one of these for a video generation task?

U Show me the status of video ID 'vid_99283'.



Checking video status... Your video 'vid_99283' is currently COMPLETED. You can watch it or download it here: [Watch Video](https://heygen.com/video/...) Should I retrieve your full video history?

U Generate a video from template 'welcome_v1' with variable {'name': 'John Doe'}.



Template generation triggered! I've started the production of your video using the 'welcome_v1' template. This usually takes 2-5 minutes. I'll monitor the status for you. Should I notify you once the URL is ready?

Frequently Asked Questions

01 How do I start using HeyGen with my AI agent?

You must first subscribe to this MCP on Vinkius and enter your specific HeyGen API Key in the agent's settings. Once that credential is stored, you can begin issuing commands like 'List all available avatars.'

02 Can I use HeyGen to translate a video?

Yes, using `translate_video` allows your agent to take an existing video and output a new version in another language while maintaining perfect lip-syncing for the speaker.

03 What is the difference between `generate_video` and `generate_from_template`?

`generate_video` creates a unique, standalone video from scratch using defined scenes. Conversely, `generate_from_template` uses pre-built templates to replace specific variables, which is better for mass personalization.

04 How do I know when my generated video is ready?

You use the `get_video_status` tool. It checks the job and will return 'completed' along with a direct URL you can immediately download or share.

05 Does HeyGen MCP handle avatar selection?







Yes, you first call `list_avatars` to see all available options. Then, you provide the specific 'avatar_id' when triggering any generation task.

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>"heygen": { "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

HeyGen 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 HeyGen. 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	HeyGen MCP
Server ID	019d75b0-d915-7371-9f80-9e7c2dd6c047
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/heygen.