MCP SERVER | NO CODE | CLOUD HOSTED

Kling AI MCP

Generate cinematic video and image assets on demand.

Kling AI (Generative Video & Image) lets you control state-of-the-art cinematic media production directly from your agent. Generate high-fidelity videos using text descriptions, animate static photos into motion, or visualize garments on models with virtual try-on. It handles everything from concept to final MP4 file.

A+ Quality Score 100/100

generative-video | text-to-video | ai-media | cinematic-generation | image-animation | creative-tools

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/kling-ai-generative-video-image

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.

Kling AI (Generative Video & Image) MCP

10 tools available
Cloud-hosted on Vinkius

Generating professional video and image assets used to be a multi-step process involving specialized software, rendering queues, and multiple export formats. Now, you talk to your AI agent like you're talking to an assistant. You give it a prompt—say, 'Show me a futuristic city at sunset.' Your agent sends that request through the MCP, managing the complex generation pipeline in the background. Whether you need a short clip from a text description or want to map a new jacket onto a model for e-commerce, your agent handles the heavy lifting. After submitting the job, it monitors the status and pulls the final high-resolution MP4 or image URLs for you to download. Connecting this capability through Vinkius means your favorite AI client can access world-class media creation without ever needing to open a dedicated studio suite.

Core Capabilities

01 — Create videos from text prompts

Submit a descriptive text prompt and receive a cinematic, high-fidelity video clip generated by the engine.

02 — Animate static images

Turn still photographs into dynamic videos by controlling the movement trajectories of the scene's elements.

03 — Visualize digital clothing on models

Blend source garment images onto target human photos to create realistic virtual try-on composites.

04 — Synchronize speech to videos

Sync audio files with a video portrait, automatically adjusting the mouth movements for professional results.

05 — Generate multiple high-quality images

Create up to four detailed images simultaneously from simple text descriptions.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/kling-ai-generative-video-image** — connect your AI agent in three steps.

01 Subscribe to this MCP and enter your dedicated Kling Access Key and Secret Key.

02 Tell your AI agent what media you need, whether it's a video or an image. The agent submits the request and gets a Task ID back.

03 Use the provided status tools with that Task ID to poll for updates until the task succeeds, then retrieve the final MP4 or image URL.

The bottom line is you talk through your preferred AI client, submit the job parameters, and let the system handle monitoring and delivery of the finished asset.

Built For

Creative Directors who need to rapid prototype storyboards; E-commerce Managers who can't wait for physical samples; or Video Editors tired of stitching together dozens of short, expensive B-roll clips.

Video Editor

Needs to quickly generate cinematic background footage (B-roll) and motion graphics using text prompts without leaving their primary editing workflow.

E-commerce Merchandiser

Must visualize new apparel collections on diverse model bodies for online listings, saving time and money compared to physical photography shoots.

Creative Director

Needs to test multiple visual concepts or storyboards by rapidly generating varied image styles or short video drafts based purely on textual prompts.

What Changes When You Connect

01 Get B-roll fast. Instead of spending hours in manual rendering software, you tell your agent to generate short, cinematic sequences using the `text_to_video` tool. You get finished footage right from the chat interface.

02 Visualize products instantly. Use `virtual_try_on` to map new clothing onto models and create realistic e-commerce assets without ever shooting a physical photo session.

03 Control every frame. If you have a static picture but need movement, use `image_to_video`. This gives your agent the power to animate photos with consistent, controlled dynamics.

04 Mass asset creation. Need multiple concepts for an ad? The `text_to_image` tool generates up to four unique images simultaneously, giving you rapid visual iteration.

05 Professional video polish. Use `lip_sync_video` when you have a speaker recording but need the mouth movements fixed. It handles the synchronization so your avatar looks natural.

Real-World Applications

Creating an ad campaign background

A brand marketer needs cinematic footage of a 'futuristic city in the rain.' Instead of hiring expensive stock footage or filming it, they ask their agent to run a `text_to_video` prompt. The agent submits the job and then uses `get_video_task` until the final MP4 link is ready for download.

Launching new seasonal apparel

An e-commerce manager needs to show how a jacket looks on five different body types. They use the agent's `virtual_try_on` tool, submitting the garment image and target model photos. The system handles all the blending, confirming success with `get_tryon_task`.

Developing explainer videos

A training department needs a spokesperson video where the speaker's mouth movements must match new audio narration. They run the `lip_sync_video` tool, and when finished, they pull the final synchronized MP4 using `get_lipsync_task`.

Building narrative storyboards

A creative director is planning a multi-scene video. They use their agent to run several varied prompts through `text_to_image`, collecting up to four high-quality images per prompt, allowing them to quickly build out a visual script.

Patterns to Avoid

Assuming immediate results

✗ AVOID

Asking the agent for a video and then immediately trying to download it without checking if the job is done. This will fail because generation takes time.

✓ INSTEAD

Always remember that media generation runs in the background. After using `text_to_video`, you must use `get_video_task` repeatedly until the status confirms success before attempting to download.

Over-relying on a single prompt

✗ AVOID

Asking for one perfect image and giving up if it's not exactly right. This wastes time when you need variety.

✓ INSTEAD

Use the `text_to_image` tool to generate multiple variations at once, since it handles up to four high-fidelity images per single request.

Ignoring job status listings

✗ AVOID

Doing a series of complex jobs and then forgetting how many were started. It's hard to track which file is where.

✓ INSTEAD

Run `list_video_tasks` or other listing tools periodically. This gives you an easy overview of all your recent video generation jobs.

The Right Fit

Use this MCP if your primary bottleneck is the *creation* and *rendering* of high-quality, complex visual media (videos, composite images). If your workflow requires text-to-video conversion, virtual try-on, or advanced image animation, this is your tool. Don't use it if you only need simple data retrieval, basic text summarization, or structured data formatting; those tasks belong with general purpose agent tools. However, if you just want to generate a single image without complex motion or high fidelity, check if an alternative

dedicated image API is faster, but for cinematic quality and variety, this MCP is unmatched.

The Pain of Manual Media Production

Today, building a single piece of marketing content requires juggling multiple specialized applications. You write the script in one place, export stills from another, and then you have to upload those assets into complex video suites—all while manually monitoring render times on separate platforms. It's slow, it’s expensive, and it involves endless copy-pasting between tools.

With this MCP, your AI agent handles the entire process conversationally. You tell it what you want, and it coordinates the generation of cinematic videos or high-fidelity images automatically. You get a finished asset link without ever having to open a dedicated rendering program.

Get Studio-Grade Visuals with Kling AI (Generative Video & Image)

Specific manual steps that disappear include the initial text-to-video prompt submission, waiting in a separate queue manager, and manually compiling the final MP4 file. Your agent takes over all those necessary status checks.

The difference is control: you tell your agent to generate multiple ideas at once using `text_to_image`, or you can run complex video processes like animating stills with `image_to_video`. You're not just getting a file; you're getting rapid, controlled creative iteration.

Kling AI (Generative Video & Image) with 10 Tools

Use these tools to manage the entire lifecycle of generative media, from creating initial concepts to retrieving final high-fidelity video and image assets.

#	TOOL	DESCRIPTION
01	text_to_video	Generates a cinematic AI video using Kling V3 from a simple text description and provides a task ID for monitoring.
02	image_to_video	Animates a static picture into a dynamic video, providing a task ID that needs to be polled until completion.
03	get_video_task	Checks the status of a video generation job and returns the final MP4 links once the process is successful.
04	list_video_tasks	Retrieves a list of your most recent Kling AI video creation jobs for tracking purposes.
05	text_to_image	Generates up to four high-fidelity images from text using the Kolors architecture, issuing a task ID for polling.
06	get_image_task	Checks the status of an image generation job and returns the final picture links when finished.
07	virtual_try_on	Maps a digital garment onto a target person's photo, providing a task ID for tracking the composite image creation.
08	get_tryon_task	Checks the status of a virtual try-on job and retrieves the final high-resolution blended image URL upon success.
09	lip_sync_video	Drives mouth movements by synchronizing specific audio files to an existing video portrait, returning a task ID.
10	get_lipsync_task	Checks the status of a lip-sync job and provides the final MP4 file when the synchronization is complete.

See It in Action

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

U Generate a 5-second cinematic video of a futuristic city in the rain using Kling V3

Video generation task submitted to Kling V3. Task ID: 'vid-98765'. I'm monitoring the render status for you. High-fidelity motion and lighting effects are being computed. I'll provide the MP4 link as soon as it's ready.

U Animate this static image into a video with gentle camera pan: [url]

Image-to-video task started. I've set the motion trajectory for a gentle horizontal pan. This will bring your static frame to life with consistent AI-generated dynamics. Your Task ID is 'ani-12345'.

U What is the status of my virtual try-on task ID 'try-456'?

Your virtual try-on task 'try-456' has SUCCEEDED. The digital garment has been naturally blended onto the model. You can view and download the final high-resolution composite here: [image-url].

Frequently Asked Questions

01 **How do I generate the best cinematic videos using text_to_video?**

Be highly descriptive. Instead of 'city in rain,' try 'cinematic 5-second shot of a futuristic neon city street slicked with rain, viewed from ground level.' The more detail you give, the better the output.

02 **What if my virtual_try_on job fails? How do I check its status?**

If it fails or is still running, use `get_tryon_task`. This tool will tell you exactly where the task is in the queue and when it's expected to complete.

03 **Can I animate an image into a video using image_to_video?**

Yes. You upload your static photo, send the request, and the system processes the motion trajectory. Once done, you use `get_video_task` to retrieve the final animated MP4.

04 **Does text_to_image generate all my required art assets?**

The `text_to_image` tool generates up to four high-quality images per request. This allows you to get multiple variations of a concept in one single call.

05 **What is the difference between text_to_video and lip_sync_video?**

Text-to-video creates an entirely new video based on words (a scene). Lip-sync video takes an existing video portrait and modifies it to match a specific audio file.

06 **How often do I need to poll for task status?**

The system documentation recommends polling every few minutes using the relevant 'get' tool (e.g., `get_image_task`). Don't spam it; wait a short period between checks.

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 `"kling-ai-generative-video-image": { "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

Kling AI (Generative Video & Image) 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 Kling AI (Generative Video & Image). 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	Kling AI (Generative Video & Image) MCP
Server ID	`019d75c1-cc32-7162-906a-a6077d98322f`
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/kling-ai-generative-video-image.