MCP SERVER | NO CODE | CLOUD HOSTED

Luma AI (Generative Video MCP)

Direct cinematic camera control and scene flow.

Luma AI Generative Video & Creative connects your agent to professional video tools. Generate cinematic videos and high-resolution images from simple text prompts or static photos, controlling camera movements like panning, tilting, and orbiting in a single conversation.

A+ Quality Score 100/100

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

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

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.

Luma AI (Generative Video & Creative) MCP

10 tools available
Cloud-hosted on Vinkius

This MCP lets you generate complex visual media using Luma's advanced generative models. You talk to your AI agent —whether it's through Claude, Cursor, or any other compatible client—and command the creation of cinematic video sequences. Need an abstract shot? Tell it. Want to animate a still photograph into something moving? Just upload the image and tell the agent what you want to see. The system handles the rendering process, giving you precise control over camera actions like dolly shots or orbits. You can also extend existing scenes or create perfect loops for background elements. Because Vinkius hosts this catalog, your agent gets access to all these professional-grade tools in one place, letting you focus only on the creative vision.

Core Capabilities

01 — Generate video from text prompts

Create high-fidelity videos from detailed descriptions of scenes or concepts.

02 — Animate still images into motion

Bring static photos to life by adding consistent, dynamic movement across the frame.

03 — Control camera movements

Specify professional shots like panning, tilting, dollying, or orbiting within your generated video.

04 — Extend and loop scenes

Continue a video past its original length or create seamless looping footage for backgrounds.

05 — Create high-res images

Generate stunning, photorealistic still images using advanced visual models.

One Click on Vinkius — From Prompt to Execution

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

01 Connect your Luma AI API Key to the Vinkius catalog and authorize access through your preferred AI client.

02 Give your agent a detailed prompt, specifying whether you want text-to-video, image animation, or camera movement parameters.

03 The MCP sends the job request. The agent tracks the status and delivers the final video link once processing is complete.

The bottom line is that you treat complex media generation like chatting with a professional editor who handles all the heavy lifting behind the scenes.

Built For

Anyone building visual content quickly. This hits the video editor drowning in manual renders, the creative director needing rapid concept mockups, and the designer whose idea needs to move off a static mood board.

Creative Director

Needs to rapidly prototype storyboards or visual treatments for clients by commanding varied camera paths and styles without waiting days for rendering.

Video Editor

Wants high-quality B-roll footage, cinematic sequences, or background elements generated on demand, eliminating the need to shoot everything manually.

Digital Artist / Designer

Needs to iterate quickly on photorealistic imagery and complex video transitions between key visual concepts directly from their development environment.

What Changes When You Connect

01 Stop manually rendering B-roll. With this MCP, you generate high-quality, cinematic sequences using text prompts or images—all within a natural conversation with your agent.

02 Gain professional motion direction. Instead of basic footage, use the `lm.camera_control` tool to precisely dictate shots, forcing movements like pans or orbits that give your content depth.

03 Improve continuity instantly. If you need a scene to continue past its initial length, use `lm.extend_video`. It seamlessly adds additional footage so your video doesn't cut off awkwardly.

04 Refine visual concepts rapidly. You can create stunning images using the `lm.text_to_image` tool and then immediately animate them into videos, speeding up your entire design loop.

05 Handle complex transitions easily. If you have two keyframes that need to connect, use `lm.interpolate`. It builds a smooth bridge between those visuals without manual frame-by-frame work.

Real-World Applications

Creating an animated product shot

A designer needs an animated visual of a new gadget for a website. They ask their agent to take a static render and use `lm.image_to_video` to bring it to life, specifying the motion should be a gentle dolly zoom over a 4-second span.

Building a cinematic background loop

A video editor requires an endless, looping shot of futuristic city streets. They prompt their agent using `lm.text_to_video` and then use the resulting job ID to initiate a status check with `lm.get_generation`, confirming the perfect loop is ready.

Storyboarding complex action scenes

A creative director wants to visualize a character moving through a canyon. They prompt their agent using `lm.camera_control` to specify a sequence that involves a slow pan, followed by an orbit shot around the main subject.

Connecting disparate visual ideas

An AI artist has two key concept images (A and B). Instead of manually blending them, they instruct their agent to use `lm.interpolate` to generate a smooth video transition that bridges the visual gap between A and B.

Patterns to Avoid

Treating it like simple stock footage

✗ AVOID

Just prompting for 'a sunny field' expecting perfect, ready-to-use video clips that match existing assets.

✓ INSTEAD

Instead of asking for generic content, define the shot entirely. Use `lm.text_to_video` with highly specific parameters like 'A wide shot of a sunlit meadow; slow pan from left to right; 5 seconds.' Always specify camera movement.

Ignoring job management

✗ AVOID

Sending the initial generation prompt and then forgetting about it, assuming the result will appear immediately.

✓ INSTEAD

Always use `lm.get_generation` with your generated ID to check if the video is 'queued,' 'dreaming,' or 'completed.' Don't assume instant results.

Trying to edit raw files outside of the agent

✗ AVOID

Downloading a video and then having to manually adjust its camera movement in separate professional software.

✓ INSTEAD

Use `lm.camera_control` directly through your agent chat interface. This embeds the specific motion parameters (pan, tilt) into the generation process itself.

The Right Fit

Use this MCP if your primary need is generating high-fidelity, cinematic *motion* and controlling camera movement within the generated media. If you require complex visual effects like blending layers or adjusting specific color grades after rendering, this isn't it—you'll need a dedicated NLE (Non-Linear Editor) tool instead. However, if your pain point is simply 'I don't have time to hire VFX

artists,' then this is perfect. Specifically, use `lm.text_to_video` when you start with nothing but an idea, or use `lm.image_to_video` when you need to animate existing photography. If the project demands perfectly structured output data (like specific metadata), check if another MCP designed for data formatting might be better suited than a media generator.

The struggle with visual pre-production today is all about manual iteration.

Right now, getting a concept visualized requires endless rounds of revisions. You write a prompt for an AI image tool, download the result, then manually take it to video software, adding motion keyframes and stitching together multiple shots just to show how the final piece would look. It's time-consuming, expensive, and involves dozens of clicks across several different programs.

With this MCP, you keep everything conversational. You tell your agent the full vision—text prompt, camera movement, even where the scene needs to extend—and it manages the entire workflow in one go. You get a ready-to-use cinematic sequence based on natural language commands.

Luma AI (Generative Video & Creative) MCP: Control motion and composition.

The biggest manual headache that vanishes is the process of generating consistent, continuous movement. You no longer have to generate a 1-second clip, then another, and manually figure out how to transition them so the camera path feels natural or if you need to force an orbit around a subject.

You just tell your agent what kind of shot you want—a 'slow dolly into the landscape'—and it handles the complex geometry and motion parameters. It’s about moving from conceptual ideas to fully realized, directed footage instantly.

Luma AI (Generative Video & Creative) - 10 Tools

Manage the entire creative pipeline, from generating initial video concepts to controlling specific camera motions and managing job status.

#	TOOL	DESCRIPTION
01	`lm.get_generation`	Checks the status of a previously submitted Luma Dream Machine job and retrieves the final video URL if ready.
02	`lm.list_generations`	Lists your most recent Luma Dream Machine creations, showing their prompts, IDs, and current status.
03	`lm.delete_generation`	Permanently deletes a specific Luma Dream Machine generation job and its associated video file.
04	`lm.extend_video`	Adds continuous footage to an existing Luma video, ensuring the scene flows naturally from the original clip.
05	`lm.interpolate`	Generates a smooth, high-quality video transition connecting two separate keyframe images using Dream Machine.
06	`lm.camera_control`	Creates videos with specific camera actions, allowing you to direct the shot with movements like pan, tilt, dolly, or orbit.
07	`lm.text_to_image`	Generates detailed, photorealistic still images based on a text description using Luma Photon-1.
08	`lm.get_credits`	Retrieves the current credit balance for Luma Dream Machine usage.
09	`lm.text_to_video`	Generates a cinematic video from scratch using only a text prompt, providing high motion coherence and photorealism.
10	`lm.image_to_video`	Animates a still image into a video by setting the photo as the initial source frame for generation.

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 spaceship landing on a neon planet

Generating cinematic video... I've triggered Luma Dream Machine with your prompt. Your Generation ID is 'gen-98765'. I'm monitoring the status. High-fidelity motion and lighting are being computed. I'll provide the MP4 link as soon as it's ready.

U Animate this static image into a video: [url]

Image-to-video task started. I'm using your image as the source frame for Luma Dream Machine. This will bring your static visual to life with consistent AI dynamics. Your Generation ID is 'gen-12345'.

U Show me my last 3 Luma generations

I've retrieved your recent generations: 1) 'Spaceship Landing' (Status: Completed), 2) 'Neon Forest' (Status: Dreaming), and 3) 'Abstract Waves' (Status: Completed). Would you like to see the video links for the completed ones?

Frequently Asked Questions

01 **How do I get started with Luma AI (Generative Video & Creative) MCP?**

You must first connect your Luma AI API Key via the Vinkius catalog. Once connected, simply instruct your agent on what kind of cinematic video you want to generate.

02 **Can I use Luma AI (Generative Video & Creative) MCP for image-to-video?**

Yes, the `lm.image_to_video` tool lets you select a static photo and animate it into video using the system's advanced motion models.

03 **What is camera control with Luma AI (Generative Video & Creative) MCP?**

The `lm.camera_control` tool allows you to generate videos that include specific, professional movements like panning across a scene or orbiting around a subject.

04 **How do I check if my video job is finished using Luma AI (Generative Video & Creative) MCP?**

Use the `lm.get_generation` tool and provide your unique Generation ID to see the current status (queued, dreaming, completed) and retrieve the final link.

05 **Is Luma AI (Generative Video & Creative) MCP better than simple GIF generation?**

Yes. This MCP generates high-fidelity video files with professional motion coherence, unlike simple looping tools that often look choppy or low quality.

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 `"luma-ai-generative-video-creative": { "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

Luma AI (Generative Video & Creative) 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 Luma AI (Generative Video & Creative). 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	Luma AI (Generative Video & Creative) MCP
Server ID	`019d75ca-f5e5-7034-ac1b-8f56c62b1a43`
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/luma-ai-generative-video-creative.