MCP SERVER | NO CODE | CLOUD HOSTED

Runway ML MCP

Generate cinematic videos from conversation.

Runway ML connects advanced video generation directly into your AI workflow. Generate cinematic clips from simple text prompts, bring still images to life with motion, or blend two keyframes for smooth transitions using models like Gen-3 Turbo and Gen-4 Turbo. It also lets you manage job status, view credits, and cancel costly renders without leaving your chat interface.

A+ Quality Score 98.33/100

video-generation | generative-ai | media-production | creative-tools | animation | model-inference

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/runway-ml

SCAN TO VIEW

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.

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.

Runway ML MCP

10 tools available
Cloud-hosted on Vinkius

Runway ML lets you generate complex video content entirely through conversation. Instead of navigating web interfaces or managing complicated pipelines, you just tell your agent what you want—a 10-second cinematic shot of a futuristic cityscape, for example. Your AI client handles the heavy lifting using best-in-class generative models. If you have an image and want to guide how it moves, you don't need complex keyframe software; you simply describe the motion. This integration puts professional video production tools right into your chat window. With Vinkius in the catalog, Runway ML sits alongside thousands of other specialized connectors, giving your agent access to a complete suite of creative and technical services. You get stunning, highly stylized videos whether you start with pure text prompts or guide animation from a single reference picture.

Core Capabilities

01 — Generate video from text

You can ask your agent to generate brand new scenes using detailed written prompts and various models.

02 — Animate still images

Bring a static picture into motion by specifying source image URLs, the model, and how long you want the final clip to be.

03 — Guide motion with text prompts

You can generate video that not only starts from an image but also uses a separate text prompt to precisely direct the movement or camera action.

04 — Smoothly transition between images

The system creates fluid, cinematic motion clips by blending two distinct keyframe images together.

05 — Manage rendering jobs

You maintain control of your workflow by listing recent tasks, checking if a render succeeded and retrieving the output link, or canceling expensive jobs to save credits.

One Click on Vinkius — From Prompt to Execution

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

01 First, activate the Runway ML integration in your AI client through Vinkius.

02 Next, sign into the developer portal and provide your secure API token for connection.

03 Finally, tell your agent exactly what you want. For instance: 'Generate a 5-second video of a cyberpunk city using Gen-3 Turbo.'

The bottom line is that once connected, your AI client treats complex video generation tasks like simple chat commands.

Built For

Creative directors and marketing teams who are tired of spending hours in dedicated rendering software or juggling multiple web dashboards. If you need to prototype sequences fast and keep the creative conversation flowing, this is for you.

Media Producer

You use this MCP to rapidly generate B-roll footage for a campaign by giving your agent specific text prompts or feeding it existing assets that need animating.

Creative Director

You prototype storyboards and conceptualize scenes within your chat, asking the agent to test different models (like Gen-4 Turbo) to see which visual direction works best before committing resources.

Social Media Content Creator

You quickly generate multiple eye-catching clips using image inputs and text prompts, dramatically speeding up the process of creating content for several platforms at once.

What Changes When You Connect

01 Rapid Prototyping: Need to test five different visual concepts? Use text_to_video or gen3_turbo to spin up quick 5-second clips, letting you rapidly iterate on creative direction without leaving your chat window.

02 Controlled Animation: Instead of guessing how an image will move, use image_text_to_video. You provide the starting photo and then direct its movement with a text prompt, giving you precise control over the animation's trajectory.

03 High Fidelity Output: For final marketing assets that need to look pristine, use gen4_turbo or text_to_video to leverage the most advanced models available for superior visual quality.

04 Seamless Blending: When your project requires a smooth shift between two distinct visuals, the interpolate tool handles it. It creates professional-grade transitional motion clips you'd otherwise need specialized software for.

05 Cost Management: Don't let renders run forever. Use list_tasks to track jobs and cancel_task immediately stops costly rendering pipelines when your creative needs change.

Real-World Applications

Developing a new product concept video

A marketing team has a mood board but no footage. They ask their agent to generate a 10-second cinematic scene of the product in action using text_to_video, and then use image_text_to_video to guide the camera movement around key features.

Creating an animated character sequence

A content creator wants a specific look for their main character. They run an initial animation using image_to_video on a reference photo, and then use interpolate to smoothly transition the character into a new setting.

Reviewing project costs

A producer finishes a day of testing videos and needs to know how many credits were used. They ask their agent to run get_organization, which immediately reports billing usage so they can plan for tomorrow's renders.

Troubleshooting failed renders

A user submits a complex prompt that fails after 15 minutes of rendering. Instead of waiting, they ask their agent to check the job status with get_task and then use cancel_task to stop it immediately.

Patterns to Avoid

Trying to render everything at once

✗ AVOID

Copying a massive prompt into an external web tool, clicking submit, and waiting for the result while doing nothing else. This process is slow, difficult to track, and requires leaving your current workspace.

✓ INSTEAD

Instead, keep it in your chat. Use text_to_video or gen4_turbo via this MCP. It sends the task to the agent, which gives you a task ID you can then poll with get_task while you continue working on other prompts.

Forgetting to check job status

✗ AVOID

Thinking that because you submitted the render request moments ago, it must be finished. You keep sending follow-up messages until your agent finally tells you it's not ready.

✓ INSTEAD

Always ask your agent to list_tasks first. This shows recent activity and confirms if the job has moved past 'PENDING' status, saving time.

Over-relying on generic image tools

✗ AVOID

Using a basic online animation tool that treats all images the same way, resulting in stiff, unnatural movements.

✓ INSTEAD

Use image_text_to_video. This specific tool allows you to combine your starting visual with a detailed text prompt, giving the AI directional instructions for realistic motion.

The Right Fit

You should use this MCP if your primary need is generating high-quality video content (cinematic shots, B-roll, animations) directly within an AI conversation. This is perfect when you are iterating quickly or managing a complex creative pipeline that involves multiple steps: text generation -> image animation -> blending ->

status checking.

Don't use this if your goal is simple file conversion (e.g., PNG to JPG). You need dedicated media converters for that. Also, don't use it just because you want video; you must have a concrete idea of the visual content first. If you only have a vague concept, write down detailed prompts using text_to_video or gen3_turbo until the vision is solid.

This MCP handles the *generation* and *management*. It's not an asset storage solution or a video editing suite; it's your engine for creating raw, high-quality visual assets.

Manual Video Generation Is a Time Sink

Today, generating even short clips means jumping between web interfaces. You upload an image to one site, write the prompt into another, and then manually manage the task ID in a third dashboard just to check if it's done. If you need to change the style or timing, you start that whole agonizing process over.

With this MCP, all those manual steps vanish. You simply describe your entire visual idea—the initial prompt, the image source, and the desired transition—and ask your agent to run the tools. What you get back is a complete video asset managed within one conversation.

Runway ML MCP Delivers Controlled Visual Assets

You no longer have to copy and paste complex instructions or manage separate dashboards. The agent handles the handoff: it takes your text input, routes it through gen4_turbo, initiates the task ID, and keeps track of its status until you ask for the final output link.

It's a fundamental shift from 'submitting a job' to 'asking the AI to produce a result.' You are in control of the entire process, from prompt to completion.

Runway ML: 10 Video Generation Tools

These tools let you run every stage of the video creation process—from initial prompt generation to final status checking and billing management.

#	TOOL	DESCRIPTION
01	`cancel_task`	Stops a video generation job that is currently running so you don't waste credits on a failed render.
02	`gen3_turbo`	Generates quick, 5-second videos using the Gen-3 Alpha Turbo model for fast testing and prototyping.
03	`gen4_turbo`	Creates high-quality video clips using the advanced Gen-4 Turbo model when you need maximum visual fidelity.
04	`get_organization`	Retrieves your account's current credit usage and overall organization details for billing tracking.
05	`get_task`	Checks the status of a specific job ID, telling you if it succeeded or failed and providing the output link when ready.
06	`image_text_to_video`	Creates video by combining a starting image with a text prompt that guides how the camera moves or what happens to the subject.
07	`image_to_video`	Animates a still picture into a moving video by specifying a source URL and desired duration.
08	`interpolate`	Creates smooth, professional-looking motion clips that transition believably between two separate keyframe images.
09	`list_tasks`	Shows a history of your recent video generation jobs so you can see what's been processed recently.
10	`text_to_video`	Generates a full video clip simply by providing a detailed text description and specifying the desired length (5 or 10 seconds).

See It in Action

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

U Create a 5 second cinematic video showing a sunset over an alien planet using Runway Gen-3 Turbo.

Your Runway generation task was successfully submitted to the Gen-3 Turbo model. The task ID is rw_task_1234. You can ask me to poll this task in a moment to review the final video output.

U Take this reference image URL and animate it with Gen-3 Turbo to make the camera slowly pan backwards.

Image to video request initiated using the `image_text_to_video` tool. Your task is now processing and I got the ID 'rw_task_9876'. The animation will be ready shortly.

U List all my ongoing tasks on Runway to see if the video has finished rendering.

I've scanned your recent 10 tasks. I see that task 'rw_task_1234' just SUCCEEDED! Here is the MP4 link to your generated video.

Frequently Asked Questions

01 **How do I use Runway ML MCP for text-to-video generation?**

Use the text_to_video tool by giving your agent a detailed description of what you want. You specify the prompt and either 5 or 10 seconds, and it returns a task ID immediately.

02 **Can I animate an existing photo using Runway ML MCP?**

Yes, use image_to_video by providing the source URL. You specify the model and duration, and the system animates that still picture into a video clip.

03 **What if my render job is taking too long? How do I stop it?**

You can use cancel_task with the specific task ID. This immediately stops the rendering process and prevents you from wasting valuable credits on an unnecessary job.

04 **How does Runway ML MCP help me blend two different videos?**

Use the interpolate tool. It takes two separate keyframe images and generates a single, smooth transitional motion clip between them for professional continuity.

05 **Where do I check my remaining credits after using Runway ML MCP?**

Run get_organization to retrieve your current organization details. This shows you exactly how many credits you've used and what your billing status is.

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 `"runway-ml": { "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

Runway ML 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 →

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by Runway ML. 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	Runway ML MCP
Server ID	`019d7601-2dba-7101-a72d-d674fa9cd4b7`
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/runway-ml.