

MCP SERVER

NO CODE

CLOUD HOSTED

# Runway ML Alternative MCP

## Automate High-End Video Generation Workflows

Runway ML Alternative automates high-end video generation directly from your AI agent. Trigger cinematic videos using Gen-3 Alpha Turbo or Gen-2 models with simple text prompts, configure precise settings like aspect ratios and duration, and monitor task progress—all without manually hitting an API endpoint.

**A+** Quality Score 100/100

video-generation

generative-ai

media-production

text-to-video

creative-automation



# 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 — 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.

**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](https://cloud.vinkius.com) — connect your AI agent in under 60 seconds.

# Runway ML MCP

4 tools available

Cloud-hosted on Vinkius

You can generate professional-grade video assets simply by talking to your AI client. Instead of navigating a complex web interface and copying manual parameters, you tell your agent exactly what you want: 'Create a cyberpunk street scene, 10 seconds long, with a 1280x768 ratio.' Your agent handles the technical calls needed for high-end video generation models. It runs the task, tracks its status in real time, and retrieves the final output URL as soon as it's finished processing. You can also check your full history of creations or delete old tasks to keep everything clean. All this capability is managed through one connection point on Vinkius, meaning you don't have to worry about which platform API key to manage; you just focus on the vision.

---

## Core Capabilities

### 01 — Generate cinematic video sequences

Create new video tasks using text prompts or by supplying an image input for models like Gen-3 Alpha Turbo.

### 03 — Track real-time task status

Query the current state of any generation job and retrieve the final output URL once processing is complete.

### 05 — Clean up old assets

Delete or cancel specific tasks from your history list to keep your work organized and clear.

### 02 — Manage advanced generation settings

Configure precise parameters such as duration (5s or 10s), specific aspect ratios, and seeds for consistent results.

### 04 — View comprehensive history

List all past video generation tasks to review assets, check metadata, or manage your workspace.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/mcp/runway-ml-alternative](https://vinkius.com/mcp/runway-ml-alternative) — connect your AI agent in three steps.

- 01 Subscribe to this MCP and provide your Runway API key.
- 02 Tell your AI agent what video you want, including details like the desired length or aspect ratio.
- 03 The agent creates the task, monitors it until completion, and gives you the final link when the video is ready.

The bottom line is that you talk to your agent like a creative director, and it handles all the complex steps of submission, waiting, monitoring, and retrieval.

---

## Built For

Creative Directors who need to rapidly prototype storyboards; Content Creators automating asset pipelines; or Developers integrating video generation into internal tools. If your job involves turning concepts into visual assets, this MCP is for you.

### Content Creator

Automates the production of large batches of video marketing assets by feeding natural language descriptions to the agent.

### Creative Director

Quickly iterates on visual concepts and storyboards, checking task statuses or listing history without leaving the chat interface.

### Developer

Integrates professional video generation capabilities directly into a coded workflow or internal application using the agent's structured outputs.

---

## What Changes When You Connect

- 01 Stop manually checking progress bars. The agent monitors the job until completion and delivers the finished output URL directly to your conversation.

- 
- 02 Generate complex video tasks (Gen-3 Alpha Turbo or Gen-2) just by describing them in plain language, eliminating the need for technical API calls.

---

  - 03 Maintain perfect organization by using `list_tasks` to review all history and `delete_task` to clear out old, unnecessary assets instantly.

---

  - 04 Gain total control over output quality by configuring specific parameters like aspect ratios, seeds, and precise durations during creation.

---

  - 05 The agent acts as a creative director for your workflow, handling the technical complexity of video generation so you can focus purely on the artistic vision.
- 

---

## Real-World Applications

### Need to prototype multiple concepts fast

A marketing team needs 10 different concept videos for a campaign launch. Instead of logging into Runway ML ten times, they ask their agent to 'create ten video tasks based on these mood boards.' The agent runs them all and lists the final URLs when ready.

### Finished with old drafts and clutter

The developer has finished a project phase and needs to clear out hundreds of draft tasks. They run `list_tasks`, identify the irrelevant task IDs, and use `delete_task` to wipe them clean, keeping their workspace manageable.

### A client review requires checking old assets

The project manager needs to show a client five videos from last month's shoot. They run `list_tasks`, see the titles, then use `get_task` on key IDs to pull up the specific finalized video links for immediate sharing.

### Integrating video into internal tools

A software developer wants a feature that generates tutorial videos based on code changes. They connect this MCP via their agent, allowing the workflow to automatically call `create_task` whenever new documentation is written.

---

# Patterns to Avoid

---

## Manual API Scripting

### X AVOID

Writing a complex script that handles task submission, polling for status changes every 30 seconds, and then extracting the final URL. This requires massive boilerplate code.

### ✓ INSTEAD

Use your AI agent to call `create_task` with all parameters in one prompt. Let the agent manage the monitoring loop internally; you just get the final result.

---

## Forgetting task IDs

### X AVOID

A user knows they need a specific video but can't remember its unique ID, so they search through complex logs manually.

### ✓ INSTEAD

Run `list_tasks` first to get the full history and identify the correct task ID. Then use that ID with `get_task` for precise status retrieval.

---

## Over-retaining drafts

### X AVOID

The workspace fills up with dozens of failed or outdated draft tasks, making it impossible to find successful jobs.

### ✓ INSTEAD

Run `list_tasks` to audit your history. Use `delete_task` on anything that is no longer needed, keeping the catalog clean and useful.

---

## The Right Fit

Use this MCP if your workflow requires automated video asset creation, status tracking, or historical management through natural language conversation. Specifically, you need to trigger models like Gen-3 Alpha Turbo using prompts, track the resulting task ID, and retrieve a final URL—that's what this does best. Don't use this if all you need is simple image generation from text; those tools handle that fine without needing complex task management. Also, don't use it if your workflow only involves modifying metadata; for pure data retrieval, another MCP might be better. But when the core action is 'Generate Video and Track It,' this is the right tool.

---

## The Pain of Manual Media Production

Think about what happens today: you write a prompt, click send. Then nothing happens. You have to open another tab and manually check if the task is running. If it's stuck, you copy the ID and paste it into a separate status checker tool. Once it succeeds, you wait for the email notification, then go back to your main document, copy the URL, and finally attach it all together.

With this MCP, your agent acts like a dedicated video pipeline manager. You give it the prompt—the full vision—and that's it. Your AI client manages the entire cycle: submission, waiting, status checking, retrieval. You get the final video link delivered back into the chat interface without ever opening another tab.

---

## Runway ML Alternative MCP: Full Video Task Management

The manual steps that disappear include checking progress bars, tracking task IDs across multiple dashboards, and manually clearing out old drafts. The agent handles these technical chores in the background.

What's different now is control. You talk to your agent like a collaborator who knows all the platform's backend processes. It's not just generating video; it's managing an entire professional production workflow.

---

# Runway ML Alternative: 4 Tools

Use these tools to create new videos, check the status of existing jobs, list your history, or delete old assets directly from your AI client.

#	TOOL	DESCRIPTION
01	<code>create_task</code>	Initiate a new video generation task in Runway ML with provided prompts and settings.
02	<code>delete_task</code>	Remove an existing, specified Runway task from your history.
03	<code>get_task</code>	Retrieve the detailed status and information for a specific video generation job.
04	<code>list_tasks</code>	Fetch a comprehensive list of all your recent video generation tasks and their statuses.

---

## See It in Action

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

**U** Create a Gen-3 Alpha Turbo video of a rainy cyberpunk street, 10 seconds long, 1280:768 ratio.



I've initiated the generation task. Your Task ID is '7a8b9c-d1e2-f3g4'. I'll monitor the status for you.

**U** Check the progress of my Runway task with ID '7a8b9c-d1e2-f3g4'.



The task is currently 'SUCCEEDED'. You can view your video here: [Output URL].

**U** List all my recent video generation tasks.



I've retrieved your task list. You have 5 recent tasks, including 'Cyberpunk Street' (Succeeded) and 'Mountain Sunset' (Succeeded).

---

## Frequently Asked Questions

### 01 How do I start a new video generation task using Runway ML Alternative MCP?

You call the `create_task` tool. You simply provide your text prompt and any necessary parameters, like desired duration or aspect ratio, through natural language to your agent.

### 02 What if I need to check the status of a video job?

Use `get_task` with the specific task ID. The MCP will tell you exactly where the job is in the pipeline—pending, running, or succeeded—and provide the final URL when done.

---

**03 Can I view all my past video generation tasks?**

Yes, run `list_tasks`. This tool pulls a comprehensive record of every task you've created, letting you see titles and status summaries at a glance.

---

**04 How do I delete old or failed Runway ML tasks?**

Use the `delete_task` tool. You just need to provide the specific ID for the task you want removed from your history, keeping your workspace clean.

---

**05 Is this MCP only good for text prompts?**

No. This MCP supports both advanced text-to-video generation and image-to-video inputs, giving you versatility in how you start a concept.







---

# Go Live in 60 Seconds

Get your connection token from [cloud.vinkius.com](https://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
 <b>Claude AI</b>	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 <b>Cursor</b>	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 <b>VS Code</b>	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"runway-ml-alternative": { "url": "..."} </code>
 <b>Windsurf</b>	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 <b>ChatGPT</b>	Settings → Tools & plugins → Add MCP server → Paste endpoint
 <b>Gemini</b>	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](https://cloud.vinkius.com) →

[vinkius.com](https://vinkius.com) · [support@vinkius.com](mailto: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 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	019e38e7-0276-7074-8027-6f253730c909
Platform	Vinkius Cloud for AI Agents
Endpoint	<a href="https://edge.vinkius.com/{token}/mcp">https://edge.vinkius.com/{token}/mcp</a>

### 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-alternative](https://vinkius.com/mcp/runway-ml-alternative).