MCP SERVER | NO CODE | CLOUD HOSTED

Cloudflare Stream MCP for AI Agents

Managing global video asset delivery and live stream preparation

Cloudflare Stream MCP connects your video infrastructure directly to any AI client. Manage complex media workflows, from initiating large uploads via TUS protocol to generating captions and configuring live simulcasts for platforms like YouTube or Twitch. You can check video metadata, list all assets, delete expired content, and even generate audio tracks—all through natural conversation.

A+ Quality Score 100/100

video-streaming | live-streaming | cloudflare | video-hosting | api-video

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

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.

Cloudflare Stream MCP

26 tools available
Cloud-hosted on Vinkius

Running a modern media operation means dealing with constant streams of video data. This MCP lets your AI agent manage that entire lifecycle without you ever having to touch the Cloudflare dashboard. You can ask it to check if an uploaded asset is ready, or prompt it to create a live input for a webinar and immediately configure simulcast outputs. Need to make sure only certain regions can view a video? The agent updates those security settings instantly. It handles everything from scheduling content deletion to initiating high-volume uploads using the TUS protocol. You'll find this connector cataloged in Vinkius, giving your AI client access to thousands of other services so you stay in the chat window and get work done.

Core Capabilities

01 — List all video assets and metadata

Lists videos in your account, allowing filtering by status (ready, error) or type (vod, live).

02 — Create a dedicated live streaming input

Creates a live input for streaming video, specifying metadata and ensuring automatic recording is enabled.

03 — Configure simulcast outputs

Generates multiple broadcast feeds from one source, setting up outputs for platforms like YouTube or Twitch.

04 — Manage content metadata & security

Updates details on a video asset, including setting allowed origins to control viewing access.

05 — Handle high-volume uploads via TUS protocol

Initiates large file uploads using the TUS protocol while attaching custom metadata and size specifications.

06 — Extract and generate media components

Generates AI captions for a video or adds secondary audio tracks via URL to enhance content.

One Click on Vinkius — From Prompt to Execution

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

01 Connect your Cloudflare Account ID and API Token within the Vinkius marketplace.

02 Your AI agent authenticates the credentials, granting it permission to manage video assets.

03 You ask your agent to perform a task—like listing all videos or creating a live input—and receive instant confirmation.

The bottom line is that you get full control over complex video workflows by talking to an interface instead of using multiple web dashboards.

Built For

This MCP is built for media operations and content engineering teams. If your job involves moving, modifying, or distributing large volumes of video assets across different platforms, this saves you hours of manual dashboard clicking.

Media Operations Manager

Manages the entire asset pipeline by listing videos and updating metadata to ensure compliance before deployment.

Live Stream Producer

Sets up live inputs and configures simulcast destinations for major platforms like Twitch or YouTube on the fly.

Developer/DevOps Engineer

Integrates video upload and management logic directly into automated scripts, using tools to initiate uploads or check storage usage.

What Changes When You Connect

01 Control the full content lifecycle from one chat window. Instead of navigating to separate upload portals, you can initiate uploads using `initiate_upload` right here.

02 Streamline compliance checks instantly. Use `get_video` or `list_videos` to pull up specific asset metadata and verify its current status without manual searching.

03 Automate broadcasting across multiple platforms. With tools like `create_live_input` and `list_live_outputs`, you can set up complex simulcasts for YouTube and Twitch in a few prompts.

04 Improve content accessibility with one prompt. You can run the `generate_caption` tool to automatically create AI captions, making videos available for wider audiences.

05 Manage security settings easily. Need to restrict viewing? Use `edit_video` to update metadata or set allowed origins instantly.

Real-World Applications

A live webinar needs a secondary feed for LinkedIn.

The agent runs `create_live_input` for the main stream, then uses `list_live_outputs` to confirm YouTube is active. Finally, it calls `create_live_output`, directing a duplicate stream specifically to LinkedIn.

A batch of videos needs watermarks and captions.

The agent first runs `generate_caption` on the whole batch. It then uses `list_watermarks` to ensure a profile exists, and finally calls `edit_video` on each asset to apply the watermark.

A developer needs to upload 50 GB of footage.

Instead of dealing with massive form submissions, the agent uses `initiate_upload`, passing necessary metadata and size constraints for the TUS protocol, handling the large data transfer in chunks.

An old video needs to be archived and removed.

The team uses `get_video` to confirm all necessary downloads are available via `list_downloads`, then confirms deletion by running `delete_video` when the asset is no longer needed.

Patterns to Avoid

Manually updating every video's security settings.

✗ AVOID

A team member logs into 20 different videos one by one to change their metadata or update the allowed origins. This takes hours and is prone to human error.

✓ INSTEAD

Ask your agent to run `edit_video` for a list of IDs, batch-updating required settings like allowed origins in a single command.

Forgetting captions on newly uploaded assets.

✗ AVOID

A video is ready, but since no one remembered the captioning step, it cannot be used by international teams or screen readers.

✓ INSTEAD

Use `generate_caption` immediately after a video is finalized. It handles the AI processing and makes the text track available for download.

Creating live streams without knowing where they'll broadcast.

✗ AVOID

A producer spins up a live input, but then realizes they also need to send it to a secondary platform like Twitch. They have to manually reconfigure the stream keys and outputs.

✓ INSTEAD

After running `create_live_input`, immediately follow up with `list_live_outputs` and use `create_live_output` to configure all required broadcast destinations at once.

The Right Fit

Use this MCP if your content workflow involves the full lifecycle of video assets, from upload (TUS protocol) through live streaming setup to final distribution. You need programmatic control over metadata updates or simulcasting outputs. However, don't use it if you only need simple file storage—a basic file transfer client is enough. If your core need is simply generating high-quality transcripts or indexes *from* the video content (without managing the asset itself), consider a dedicated transcription service instead of relying on `generate_caption` alone.

Cloudflare Stream MCP: Managing Video Asset Metadata and Security

Right now, changing basic video details—like adding a specific legal disclaimer to the metadata or restricting who can view it—requires jumping into multiple tabs and clicking through complex forms. You have to manually update allowed origins for every asset.

With this MCP, you simply ask your agent to run `edit_video` on an ID list, specifying the new metadata or security rules. It handles the entire sequence of updates in a single operation. The result is immediate, auditable changes across your video library.

Cloudflare Stream MCP: Automating Live Broadcast Outputs

Setting up live streaming used to be a multi-step nightmare: create the input, confirm the stream key, and then manually configure every single output destination (YouTube, Twitch, etc.). Any missed step meant lost broadcast time.

Now, you run `create_live_input` for your main feed. Then, with one follow-up prompt, you can call `list_live_outputs` and use `create_live_output` to define every simulcast destination instantly. You're managing an entire global broadcast chain in chat.

26 Tools for Video Asset & Live Stream Management

Orchestrate every stage of your media pipeline, from initiating large video uploads to configuring complex simulcast outputs using these specialized tools.

#	TOOL	DESCRIPTION
01	copy_audio_track	Adds a secondary audio track to an existing video asset using a provided URL.
02	create_live_input	Creates and configures a new live input stream with specified recording options for Cloudflare streaming.
03	create_live_output	Sets up a simulcast output, routing the feed to a specific platform like YouTube or Twitch.
04	create_m4a_download	Enables direct M4A audio downloads for a particular video asset.
05	create_mp4_download	Enables direct MP4 video downloads for a specific video asset.
06	create_signing_key	Generates and manages secure Stream Signing Keys needed for restricted content access.
07	create_watermark	Creates a defined watermark profile to brand the video content.
08	delete_live_input	Permanently removes an active or scheduled live input stream configuration.
09	delete_live_output	Deletes a specific simulcast output profile that was configured for a live stream.
10	delete_video	Permanently removes an entire video asset and all associated copies from the account.
11	edit_video	Updates general metadata or changes settings for an existing video file.
12	generate_caption	Uses AI to automatically create accurate text captions (subtitles) for a video asset.
13	get_storage_usage	Retrieves the account-wide statistics detailing current storage consumption.
14	get_video	Fetches all specific details, status, and size information for a single video ID.
15	get_webhook	Views the current configuration settings for automated webhooks.
16	initiate_upload	Starts a large video upload session using TUS protocol, requiring metadata and size input.

#	TOOL	DESCRIPTION
17	`list_audio_tracks`	Lists all additional audio tracks currently associated with a specific video asset.
18	`list_captions`	Retrieves a list of available text track files, such as subtitles or captions, for a video.
19	`list_downloads`	Shows all potential download formats and links available for a given video.
20	`list_live_inputs`	Retrieves a list of all currently configured live input streams.
21	`list_live_outputs`	Lists all simulcast outputs available for a specific, active live input stream.
22	`list_signing_keys`	Displays all existing Stream Signing Keys used across the account.
23	`list_videos`	Lists every video in your Cloudflare Stream account, allowing filtering by status or type.
24	`list_watermarks`	Displays all created watermark profiles used for branding content.
25	`update_webhook`	Creates or modifies the webhook configuration to alert external systems about video events.
26	`upload_caption`	Uploads a new text track file, such as a VTT caption file, for synchronization with a video.

See It in Action

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

U I need a list of all my video assets, but only the ones that are marked as 'error' or 'in-progress'.

Video Asset Status Report

- `vid-921x` : Failed to process due to codec mismatch (Error)
- `vid-45b2` : Currently processing, ETA 30 mins (In Progress)
- `vid-11f7` : Processing... (In Progress)

Please check the webhook logs for details on `vid-921x`. Would you like me to try re-triggering the error video?

U Can we set up a live stream called 'Q3 Town Hall' and make sure it goes to YouTube, but only records automatically?

Live Input Setup Complete

✅ Live Input: Q3 Town Hall
- RTMPS URL: rtmp://stream.example.com/live/q3-townhall
- Stream Key: ABCXYZ123
- Recording: Automatic (Enabled)

I also configured the simulcast output for YouTube. You're good to go!

U What are the download options and formats available for the 'Product Demo 2024' video?

Downloads Available for Product Demo 2024

FORMAT	TYPE	RESOLUTION	ACTION
MP4	Video	1080p	[Download Link]
M4A	Audio	Stereo	[Download Link]
WebM	Video	720p	[Download Link]

Note: All links expire in 7 days. We can also generate a custom signing key if you need restricted access.

Frequently Asked Questions

01 **How do I manage and update video metadata using the Cloudflare Stream MCP for AI Agents?**

You tell your agent exactly what changes you want. Instead of manually logging into a dashboard, you just ask it to run `edit_video` with the new settings or origins. This instantly updates the asset's profile.

02 **Can I use Cloudflare Stream MCP for AI Agents to set up simulcasting?**

Yes, absolutely. You can tell your agent to create a live input and then specify all necessary outputs like Twitch or YouTube using the dedicated tools. It handles connecting multiple platforms in one workflow.

03 **What if I need to upload a massive amount of footage? Does Cloudflare Stream MCP for AI Agents handle it?**

It does, via the TUS protocol. You initiate the upload using `initiate_upload` and provide metadata. The agent manages the large file transfer process in secure chunks automatically.

04 **Does Cloudflare Stream MCP for AI Agents help with accessibility features like captions?**

Yes, it does. You can trigger `generate_caption` on any video asset using a simple prompt. The agent runs the AI processing and makes the resulting caption file immediately available.

05 **I need to check my storage usage for my videos. How do I use Cloudflare Stream MCP for AI Agents?**

Just ask your agent to retrieve the account-wide statistics using `get_storage_usage`. It gives you a clear, single number showing how much space your video library is consuming.

06 **How can I permanently delete old videos or live inputs with Cloudflare Stream MCP for AI Agents?**

You simply ask the agent to run `delete_video` on the asset ID, or use `delete_live_input` if you're done with a broadcast. The process is instant and irreversible.

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 `"cloudflare-stream": { "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

Cloudflare Stream 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 Cloudflare Stream. 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	Cloudflare Stream MCP
Server ID	`019e3879-1476-7336-85af-8d3537e27bdd`
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/cloudflare-stream.