MCP SERVER | NO CODE | CLOUD HOSTED

Tencent Cloud VOD MCP

Manage Assets and Audit Storage via Natural Conversation

Tencent Cloud VOD MCP connects your AI client to Tencent's video-on-demand infrastructure for deep asset management, task tracking, and storage auditing. Your agent can instantly search massive media libraries by keywords or tags, pull detailed metadata on specific files, manage content classification, and monitor background processing tasks in real time. It gives you a single conversation point to handle everything from updating file descriptions to checking overall storage volume statistics without needing the complex console.

A+ Quality Score 100/100

video-on-demand | media-processing | asset-management | video-metadata | storage-auditing | transcoding

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/tencent-cloud-vod

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.

Tencent Cloud VOD / 腾讯云点播 MCP

8 tools available
Cloud-hosted on Vinkius

Managing video assets used to mean navigating dense dashboards and running multiple reports manually. Now, your agent handles it all via this MCP. You can ask your AI client to perform sophisticated actions on Tencent Cloud VOD. Need to know the technical metadata for a batch of files? Just ask. Want to find every piece of media related to 'Q4 Marketing'? Your agent searches the entire library automatically. It also monitors complex background jobs, giving you real-time status updates on transcoding tasks or helping audit your storage usage over specific time periods. Because this MCP is hosted on Vinkius, you connect once and gain immediate access to this specialized VOD intelligence alongside thousands of other services from any compatible AI client.

Core Capabilities

01 — Audit Storage Usage

Retrieves detailed statistics about your overall storage volume for specific time frames.

02 — Search Media Libraries

Searches the entire VOD catalog using advanced filters like keywords, tags, or media classes.

03 — Manage File Details

Reads specific metadata about a video file or updates its title and description properties.

04 — Track Background Tasks

Checks the current status and technical details of asynchronous processing jobs, like transcoding.

05 — Delete Assets

Removes specific media files from your VOD infrastructure to clean up obsolete content.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/tencent-cloud-vod** — connect your AI agent in three steps.

01 Subscribe to this MCP and provide your Tencent Cloud SecretId, SecretKey, and optional SubAppId credentials.

02 Connect the service using your preferred AI client (Claude, Cursor, etc.) via the Vinkius platform.

03 Use natural language prompts in your client—your agent performs complex operations like searching or auditing storage directly.

The bottom line is you tell your agent what needs doing, and it handles all the API calls to Tencent Cloud for you.

Built For

This MCP is essential for video operations teams, content managers handling massive archives, and developers integrating VOD features. It solves the pain point of constantly jumping between dashboards just to track a file's status or audit storage limits.

Video Operations Engineer

They monitor processing health across thousands of assets, using this MCP to check `get_task_detail` for failed transcodes and immediately initiate cleanup via `delete_media_file`.

Content Manager

They need to quickly find and verify specific content. They use the agent to run a targeted search with `search_media_files`, filtering by class or tag to narrow down huge media collections.

DevOps Engineer

They integrate VOD features into applications, using this MCP's tools like `get_storage_stats` to write code that respects cloud usage limits and performs accurate capacity planning.

What Changes When You Connect

01 Instead of manually running searches, your agent uses `search_media_files` to instantly find assets across your entire library using keywords or tags. It’s immediate content discovery.

02 You don't have to log into the console just to check stats. Use `get_storage_stats` and ask your client for a usage report covering any specific time range, giving you clear capacity data.

03 When content needs cleanup, use `delete_media_file`. You can authorize the removal of obsolete assets without manually verifying every single ID number.

04 Keep track of complex jobs using `get_task_detail`. Your agent monitors background processing tasks and reports back on success status or errors instantly.

05 Need to fix metadata? Use `update_media_info` to change titles or descriptions for multiple files, all through a simple command. No bulk CSV edits required.

Real-World Applications

Auditing an old campaign's content

A content manager needs to know if any videos from the 'Summer Sale 2023' campaign still exist. Instead of manually browsing, they ask their agent to run `search_media_files` using the keyword and class tag. The agent returns a definitive list of all remaining files.

Verifying post-transcode assets

A video engineer finishes uploading 50 videos that need transcoding. They ask their agent to run `get_task_detail` for the batch job ID. The agent confirms if all tasks are 'FINISH' and ready for use.

Cleaning up expired content

The team discovers a folder of promotional videos that haven't been touched in two years. They instruct their agent to run `list_recent_media` first, confirm the files are obsolete, and then use `delete_media_file` on the confirmed list.

Capacity planning review

The CTO wants a report showing total cloud storage usage for Q1. They prompt their agent to execute `get_storage_stats`, receiving precise data points needed for budgeting without leaving the chat interface.

Patterns to Avoid

Trying to update files one by one

✗ AVOID

The user finds 15 videos that need a new description. They manually copy and paste the same updated text into the web console for each file, wasting time.

✓ INSTEAD

Use the agent to run `get_media_info` first to gather all necessary FileIds, then use the combined list with `update_media_info` in one go. Your AI client handles the batch update logic.

Confusing search and listing

✗ AVOID

The user asks their agent to 'show me everything', which results in a massive, unusable dump of every file ever uploaded.

✓ INSTEAD

Always use `search_media_files`. Specify the criteria (like tags or classes) so the agent narrows the scope. For example: 'Search for files tagged with #ProductLaunch'.

Assuming a task is complete

✗ AVOID

A video goes missing after upload. The user assumes it failed, but they don't know if the transcode job even ran or what its status was.

✓ INSTEAD

Always check the background process first. Use `get_task_detail` with the specific task ID to confirm the actual processing state before diagnosing content issues.

The Right Fit

Use this MCP if your primary need is controlling and auditing assets *within* a highly specialized, existing VOD platform like Tencent Cloud. This tool excels when you need granular control over metadata (`get_media_info` , `update_media_info`), deep search capabilities (`search_media_files`), or detailed resource tracking (`get_storage_stats`). Don't use this if your problem is general media storage across multiple, disparate cloud services; in that case, look for a generic Cloud Storage auditing tool. Also, don't rely on it to *create* new content—it only manages existing assets and processes. If you need workflow orchestration involving billing or user management outside of the video platform itself, another specialized service is better.

The Headache of Video Asset Management

Today, managing a large media library means logging into a console and clicking through tabs. You run one report to check storage, then open another screen to see the video metadata, and finally run a third job just to confirm if a background transcode finished successfully. It's constant context switching and painful copy-pasting of IDs.

With this MCP, you talk to your agent like you're talking to a coworker who knows everything about VOD. You ask the question—whether it's 'What was our peak storage usage last quarter?' or 'Find me all videos marked for archival.'—and get a direct answer without ever clicking an unnecessary button.

Tencent Cloud VOD / 腾讯云点播: Full Asset Control

The manual process of checking status, updating descriptions, and auditing storage is removed. You no longer need to run separate queries for different tasks; the agent ties it all together.

It's simple: you state the operational goal, and your AI client executes the complex series of calls—from `search_media_files` to `get_task_detail` —providing one definitive result.

Tencent Cloud VOD / 腾讯云点播 with 8 Tools

Use these eight specific tools to perform detailed operations on your VOD library, including searching files, modifying metadata, and checking storage usage.

#	TOOL	DESCRIPTION
01	`delete_media_file`	Permanently removes a specified media file from the VOD platform.
02	`get_media_info`	Retrieves detailed technical metadata for any given video asset.
03	`get_storage_stats`	Pulls a usage report detailing your total storage volume over a defined time period.
04	`get_task_detail`	Checks the real-time status and technical output of an asynchronous VOD task.
05	`list_media_classes`	Displays a list of all media content categories or classes used in your project.
06	`list_recent_media`	Fetches a list of media files that have been recently updated or modified.
07	`search_media_files`	Searches the entire VOD library using keywords, tags, and advanced filters.
08	`update_media_info`	Allows you to modify basic properties of a media file, like its title or description.

See It in Action

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

U Search for media files containing 'Product Launch' in my VOD library.

I've searched your VOD library. I found 3 files matching 'Product Launch', including '2023_Product_Launch_HQ' (FileId: `88210934`). Would you like me to retrieve the full technical metadata for these files?

U Check the status of VOD task 'task-1234567890'.

The VOD task `task-1234567890` is currently 'FINISH'. It was a Transcode task that completed successfully at 14:30. Should I check if the resulting media files are now indexed in your library?

U Show me our VOD storage usage stats for the last 30 days.

I've retrieved your storage usage data. Over the last 30 days, your peak storage reached 1.2 TB, with an average daily increase of 5 GB. Your current active SubAppId is set to 'Root Application'. Would you like a daily breakdown of these metrics?

Frequently Asked Questions

01 **How do I check storage usage using Tencent Cloud VOD / 腾讯云点播?**

You use the MCP's tool, `get_storage_stats`. Simply ask your agent to run this function and specify the exact time range you need data for. It pulls a detailed report directly.

02 **Can I update video titles using Tencent Cloud VOD / 腾讯云点播?**

Yes, that's handled by `update_media_info`. You tell the agent which files need changes, and it modifies the specified properties like title or description for you.

03 **Is Tencent Cloud VOD / 腾讯云点播 good for finding old media?**

It's great. Use `search_media_files` to query your library. You can filter by tags, keywords, and classes to pinpoint assets quickly without browsing.

04 **What is the difference between searching and listing files with Tencent Cloud VOD / 腾讯云点播?**

Use `list_recent_media` when you just want a general view of what was uploaded or changed lately. Use `search_media_files` when you need to filter by specific criteria like 'Product Launch' tags.

05 **Can I delete videos with Tencent Cloud VOD / 腾讯云点播?**

Yes, the tool `delete_media_file` lets your agent remove assets. Be sure you confirm the FileId before letting it run to prevent accidental data loss.

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 `"tencent-cloud-vod": { "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

Tencent Cloud VOD / 腾讯云点播 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 Tencent Cloud VOD / 腾讯云点播. 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	Tencent Cloud VOD / 腾讯云点播 MCP
Server ID	`019d848b-4814-7200-974e-e94da9b5a73d`
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/tencent-cloud-vod.