MCP SERVER | NO CODE | CLOUD HOSTED

Sonix MCP

Transcribe, Summarize, and Translate Media Content

Sonix lets you turn raw audio and video files into structured data, summaries, and translations using natural conversation. Transcribe full transcripts in plain text or SRT format, automatically summarize large batches of recordings, and prepare media for global audiences—all without leaving your AI client.

F Quality Score 3.6/100

transcription | subtitles | translation | media-management | ai-summary

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/sonix

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.

Sonix MCP

30 tools available
Cloud-hosted on Vinkius

Use this MCP to handle all parts of media post-production directly from your agent. Whether you're working with hours of interviews or dozens of podcast clips, you don't have to manually upload them into a separate web portal. You can tell your AI client to transcribe the raw audio and get plain text transcripts instantly. Need subtitles? It handles that too. If you need global reach, it translates those transcripts into multiple languages automatically. Furthermore, if you have a large folder of recordings, you can request a batch summarization for every file at once. This makes your agent act like a true media assistant, handling everything from generating specific video formats to managing user access and organizing the library in folders.

Core Capabilities

01 — Transcribe audio/video into text

The MCP generates transcripts in plain text, SRT, VTT, or JSON formats with precise speaker labels and timestamps.

02 — Summarize media content

You can create a summary for a single file or run batch summarization across an entire folder of recordings.

03 — Translate transcripts to new languages

Automatically process and translate existing text transcripts into dozens of different languages.

04 — Organize media library structure

The agent can list, create, update, or delete folders and individual media files within your Sonix account.

05 — Prepare video for social sharing

Initiate processes to burn subtitles directly onto the video track, creating content ready for immediate upload.

06 — Manage team access and permissions

You can list current users, invite new team members, or generate secure share links for specific media files.

One Click on Vinkius — From Prompt to Execution

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

01 First, connect your Sonix API key to the Vinkius Catalog using your preferred AI client.

02 Next, give your agent a direct command, like 'Transcribe this video and summarize it,' referencing the media file's ID or location.

03 Finally, the MCP executes the task through its tools, returning status updates and processed data—whether that's plain text, an SRT file, or a summary report.

The bottom line is you talk to your agent like it’s already connected to your media backend, bypassing manual web portal steps entirely.

Built For

Any role that deals with large volumes of spoken word or video content needs this. If your job involves converting raw footage into usable, searchable text, you need Sonix. It cuts out the manual work of exporting data and copy-pasting between services.

Content Creator

A Content Creator uses this to take a finished podcast episode, ask their agent for an instant summary, generate subtitles for YouTube, and translate the transcript for international promotion.

Journalist / Researcher

They use this MCP to quickly process hours of interview footage. Instead of watching everything, they ask the agent to extract text transcripts, then search those texts for specific keywords or names across multiple files.

Product Marketing Manager

A PM uses this when analyzing customer feedback calls. They feed the recordings into the MCP and use it to generate summaries of key pain points or feature requests, which they then share with engineering via automated links.

What Changes When You Connect

01 Instant transcripts in multiple formats: You get the raw text (using `get_transcript_text`), or structured files like SRT/VTT for direct use in video editing software. No manual conversion needed.

02 Efficiency through batch processing: Instead of running summarization on 20 clips one by one, you can initiate a whole folder summary using `create_batch_summarization` and check status with `get_batch_summarization`.

03 Global reach from local files: Need to hit multiple markets? You ask the agent to run a translation via `create_translation`, giving your content instant multilingual visibility for global campaigns.

04 Full media lifecycle control: From listing all available files (`list_media`) to organizing them into project folders (`create_folder`), you manage the entire asset pipeline without leaving your AI client.

05 Streamlined collaboration: You can invite team members using `invite_user` and manage access by generating specific, trackable share links with `create_share`, keeping your media organized and secure.

Real-World Applications

Analyzing an Interview Series

A journalist has 10 hours of raw interview footage. They ask their agent to list all the files (`list_media`), then run `create_summarization` on each one, and finally compile a master document using the summary reports. This saves days of manual reading.

Preparing Content for YouTube Launch

A content creator finishes an episode. They ask their agent to transcribe it (`get_transcript_text`), then run `create_video_burn_in` so the subtitles are baked into the video, and finally use `create_translation` to get Spanish versions for a dual-market launch.

Onboarding New Team Members

A marketing manager needs to give access to three specific folders of brand assets. They ask their agent to list existing users (`list_users`), then create the necessary folders, and finally use `create_share` to grant temporary viewing rights only.

Reviewing User Feedback Calls

A product manager needs insights from a week's worth of recordings. They ask their agent to process all files in the folder using `create_batch_summarization` and then retrieve detailed, word-level timestamps for key moments using `get_transcript_json`.

Patterns to Avoid

Manually downloading transcripts

✗ AVOID

Recording a call, logging into the Sonix website, clicking 'Download Transcript', selecting format (SRT), and then having to manually upload that file elsewhere. This is slow.

✓ INSTEAD

Just tell your agent, 'Get the transcript for this media ID in SRT format.' The MCP handles all the downloading and formatting instantly using `get_transcript_srt`.

Handling multiple languages

✗ AVOID

Transcribing a video to English, then manually hiring a translator, getting them to deliver a separate file, and uploading it. This adds huge friction.

✓ INSTEAD

Run the transcription first, then immediately tell your agent to run `create_translation` for all required languages in one go.

Fragmented project management

✗ AVOID

Losing track of where assets are. You have files scattered across different folders and don't know which ones are ready to share.

✓ INSTEAD

Use `list_media` to see everything, then use `create_folder` to create a dedicated 'Project X Assets' folder, and move them all there via the agent.

The Right Fit

You need this MCP if your workflow revolves around taking spoken media—audio or video—and turning it into structured, actionable data. If you consistently find yourself downloading transcripts, summarizing clips, or translating content across multiple platforms, this is essential. Don't use this if your primary goal is just simple file storage; for that, basic cloud connectors work fine. However, if you need to *process* the contents of the files (i.e., get text out of them),

you must use Sonix. For example, don't just list media files using `list_media` —if you want those files summarized or translated, you have to initiate one of the processing tools like `create_summarization`. It's a content processor, not just an organizer.

The Media Workflow Mess

Right now, managing media is clicking through five different tabs: the upload portal, the transcription dashboard, the subtitle editor, the translation service, and then a project management board. You download an SRT file here, copy-paste the summary into Notion there, and manually track who saw which share link.

With this MCP, your agent acts as the central hub. You simply tell it to get the full text transcript for all files in a folder, create video burn-in subtitles, translate them to French, and organize everything under a new 'Europe Launch' folder. Everything happens through natural conversation.

Sonix: Structured Data from Raw Footage

You don't have to waste time exporting data into formats that require further cleanup or re-uploading. The MCP handles generating the plain text transcript, getting word-level timestamps using `get_transcript_json`, and even updating speaker labels automatically.

What changes is the friction point. You stop managing file transfers and start directing intelligence. Your agent delivers ready-to-use, structured data directly into your workflow.

Sonix: 20 Media Processing Tools

These tools let you programmatically manage every stage of the media lifecycle—from listing files to generating translations and summaries.

#	TOOL	DESCRIPTION
01	`create_batch_summarization`	Generates a summary report for an entire folder containing multiple media files.
02	`create_folder`	Creates a new, organized container (a folder) within your Sonix media library.
03	`create_media_export`	Initiates the creation of a downloadable package containing one or more media files.
04	`create_share`	Generates a specific link that allows another user to view a selected media file.
05	`create_summarization`	Creates a summary report specifically for one single media file.
06	`create_translation`	Starts the process of translating an existing media transcript into a different language.
07	`create_video_burn_in`	Prepares a video by permanently overlaying subtitles onto the video track, making it ready for social media use.
08	`delete_media`	Removes a specific media file from your Sonix account library.
09	`delete_share`	Revokes access by removing an existing share link for a given media file.
10	`get_batch_summarization`	Retrieves the status and details of a previously requested batch summary job.
11	`get_media_export`	Checks the current progress or completion status of a media export request.
12	`get_media`	Gets general details and status information for any piece of media in your account.
13	`get_summarization`	Retrieves the final summary text or current processing status for a single file's summarization job.
14	`get_transcript_json`	Fetches a detailed transcript that includes timestamps linked to specific words spoken in the audio.

#	TOOL	DESCRIPTION
15	`get_transcript_srt`	Downloads the media's transcript formatted as an industry-standard SRT file, useful for video editors.
16	`get_transcript_text`	Retrieves a clean, continuous text dump of the entire audio content without time stamps or formatting.
17	`get_transcript_vtt`	Downloads the media's transcript formatted as a VTT file, common for web video players.
18	`get_translation`	Checks the status and retrieves the translated content from a previous translation request.
19	`get_video_burn_in`	Checks the progress of creating burn-in subtitles on a video file.
20	`invite_user`	Invite a new user to the account
21	`list_folders`	List all folders
22	`list_media`	List media files
23	`list_shares`	List shares for a media file
24	`list_users`	Lists all user accounts that currently have access to your media library.
25	`split_transcript`	Automatically split transcript into subtitles
26	`submit_media`	Submit new media for transcription
27	`update_folder`	Update a folder
28	`update_media`	Update media attributes
29	`update_transcript`	Update transcript words and speakers
30	`update_user`	Update a user role

See It in Action

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

U List my most recent media files in Sonix.

I've retrieved your media list. You have 5 recent files, including 'Interview_v1.mp4' (ID: m123) and 'Keynote_Speech.mp3' (ID: m456). Would you like to get the transcript for any of these?

U Transcribe this audio file: https://example.com/audio.mp3 in French.

I've submitted the media for transcription. The Media ID is 'm789' and the status is currently 'processing'. I can notify you once the French transcript is ready.

U Generate a summary for media ID m123.

I've initiated the summarization for 'Interview_v1.mp4'. You can check the result using `get_summarization` with ID 's999' in a few moments.

Frequently Asked Questions

01 **How do I get a plain text transcript using Sonix MCP?**

You request the transcript type 'text' from your agent. This provides a clean, continuous dump of all spoken words without any time stamps or formatting.

02 **Can Sonix MCP summarize multiple videos at once?**

Yes, you can use `create_batch_summarization`. You point it to an entire folder, and the MCP handles running the summary job on every file within that container.

03 **What format does Sonix provide for subtitles?**

It provides several formats. For professional video editing, you can use `get_transcript_srt` (SRT). For web display, the VTT format is available via `get_transcript_vtt`.

04 **How does Sonix MCP help with team access?**

You manage access by listing current users using `list_users`, or you can invite new members and generate secure share links for specific media assets.

05 **Is the translation from Sonix MCP automatic?**

Yes, after transcribing a file, you use `create_translation` to automatically process the text into dozens of required languages without human intervention.

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 `"sonix": { "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

Sonix 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 Sonix. 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	Sonix MCP
Server ID	`019e5d57-328b-718e-97df-825e8bde2331`
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/sonix.