MCP SERVER | NO CODE | CLOUD HOSTED

AssemblyAI MCP for AI Agents

Process and Analyze Spoken Word from Audio and Video Files

AssemblyAI lets your AI client transcribe audio and video files with extreme accuracy, finding more than just words. It automatically identifies who is speaking, summarizes content, analyzes mood, and even chapters out long recordings so you can process complex media instantly.

A+ Quality Score 100/100

speech-to-text | transcription | audio-intelligence | natural-language-processing | speech-recognition | video-analysis

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

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.

AssemblyAI MCP

9 tools available
Cloud-hosted on Vinkius

Stop manually uploading files to web portals or waiting for slow human transcription services. This MCP lets your AI agent take full control of high-fidelity audio intelligence right inside your workflow. You point it at a public video URL, and it handles the heavy lifting.

Your agent can transcribe speech using advanced models that deliver superhuman accuracy. Beyond just text, it automatically figures out who said what by identifying individual speakers. It also pulls out deep insights like automated summaries, topic breakdowns, or even sentiment—telling you if the discussion was positive or negative at specific points in time. When connected via Vinkius, your AI client acts as a dedicated audio engineer and linguistic analyst, making content discovery simple enough to manage right from your conversation.

Core Capabilities

01 — Transcribe Audio/Video URLs

Sends an external link to the MCP and receives a highly accurate transcript of all spoken content.

02 — Determine Speakers and Dialogue

Separates the transcription into distinct segments, labeling exactly which speaker spoke at any given moment.

03 — Generate Automated Summaries

Creates concise summaries of long recordings, giving you the key takeaways without reading through every word.

04 — Analyze Sentiment and Topics

Pulls out high-level insights by detecting overall mood (sentiment) or specific themes (topics) within the speech.

05 — Map Content Chapters

Creates an automated chapter breakdown of media, helping you navigate long videos or podcasts instantly.

One Click on Vinkius — From Prompt to Execution

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

01 You subscribe to this MCP on Vinkius and retrieve your API key.

02 Your AI client sends the external audio or video URL and specifies what insight you need (e.g., 'Summarize and detect sentiment').

03 The MCP processes the data, returning the requested structured insights—be it a transcript, summary, or speaker list—directly to your agent.

The bottom line is that you just pass an audio link to your AI client, and all the intelligence comes back formatted for immediate use.

Built For

This MCP is essential for anyone drowning in media content. It helps developers building complex pipelines, support teams summarizing customer calls, or content creators who need to quickly turn raw video into publishable text.

Content Creator

Needs to instantly generate podcast transcripts and chapter markers from recorded interviews so they can prepare show notes without spending hours transcribing.

Support Team Manager

Must summarize long customer service calls and analyze the overall sentiment across hundreds of recordings to spot emerging product pain points.

Developer

Needs to integrate high-speed speech-to-text intelligence into custom business workflows, using the transcripts to populate databases or trigger automated alerts.

What Changes When You Connect

01 Automate content discovery: Instead of manually reading minutes of meeting notes, you can use the `get_summary` tool to get immediate, high-fidelity executive reports.

02 Identify speaker contributions: The ability to label speakers via `get_speakers` ensures that every person's contribution is perfectly documented for meeting minutes and interviews.

03 Mine emotional intelligence: Use `get_sentiments` to analyze customer feedback or sales calls, pinpointing exactly when the mood shifted from positive to concerned.

04 Structure your media library: The automated chapters provided by `get_chapters` let you treat massive video files like perfectly indexed articles, improving content accessibility.

05 Full control over history: You can use `list_transcripts` and then manage specific jobs with `delete_transcript`, keeping your records clean and organized.

Real-World Applications

Analyzing Customer Support Calls

A support manager asks their agent to process a week's worth of call recordings. The agent uses the MCP to automatically generate summaries, analyze sentiment for high-risk calls using `get_sentiments`, and pull out key topics via `get_topics` to identify training needs.

Creating Podcast Show Notes

A content creator uploads a raw interview video URL. The agent first runs the transcription (`transcribe_audio_url`), then uses `get_speakers` and `get_chapters` to build detailed, multi-section show notes ready for immediate publishing.

Indexing Legal Interviews

A paralegal needs quick insights from dozens of recorded interviews. They run the MCP to get speaker labels (`get_speakers`) and use `get_summary` on each transcript, ensuring every key point is captured for litigation support.

Monitoring Webinars

A sales team needs a quick recap of a webinar. The agent runs the MCP to detect topics (`get_topics`) and get an overall sentiment report on the Q&A session, providing immediate follow-up talking points.

Patterns to Avoid

Manual Transcription Uploads

✗ AVOID

Copying a link from a video site into a third-party web portal and waiting hours for the service to process it, then manually downloading multiple files.

✓ INSTEAD

Instead, pass the URL directly through your AI client. Use `transcribe_audio_url` with your agent. This keeps the whole workflow conversational and requires no manual file handling.

Generic Summary Tools

✗ AVOID

Getting a vague summary that just restates facts without telling you *why* those facts matter or who said them.

✓ INSTEAD

Use `get_summary` combined with the speaker identification tool, `get_speakers`. This gives you an accountable summary: 'John suggested X, and Jane disagreed because Y.'

Ignoring Contextual Data

✗ AVOID

Running a transcription but failing to analyze the emotional state of the conversation, missing key points about customer frustration.

✓ INSTEAD

Always pair `transcribe_audio_url` with `get_sentiments`. This adds necessary context, letting you know not just *what* was said, but *how* it was said.

The Right Fit

Use this MCP if your primary pain point is turning raw audio or video into structured, analyzable data. Specifically, if you need to know who spoke (use `get_speakers`), what the mood was (check `get_sentiments`), or how long the content naturally breaks down into sections (look at `get_chapters`).

Don't use this if you only need simple text extraction from a document PDF. For that, a general OCR tool is better. Also, if your data source requires physical file uploads to a private server (not a public URL), check the documentation first. This MCP works best with easily accessible links.

If you are building complex pipelines, remember you can combine tools; for example, transcribe an audio link and then immediately ask

your agent to run `get_topics` on the result. It's designed for chaining intelligence together.

AssemblyAI MCP: Streamlining Media Analysis with Speech Intelligence

Today, analyzing recorded content means a lot of copy-pasting. You grab a meeting recording link, upload it to one service for transcription, download the text, then upload that text to another tool just to get a summary. It's a multi-step chore involving multiple dashboards and manual handoffs.

With this MCP, your AI agent handles the entire sequence in plain conversation. You pass the URL once, and the agent manages the full lifecycle: it transcribes the speech, then automatically runs `get_summary` and `get_speakers`. The result is not just a file; it's an actionable report handed directly back to you.

AssemblyAI MCP: Tracking Content History Using Transcription Management

Before this, keeping track of your transcripts was messy. You had files scattered across different cloud folders, and if you needed the status of a job you started last week, you either waited or manually checked multiple dashboards.

Now, your agent maintains a clean record for you. Use `list_transcripts` to see everything you've done, check the status with `get_transcript`, and even clear out old data using `delete_transcript`. It keeps your entire media archive organized in one place.

9 Tools in the AssemblyAI MCP for Audio Transcription & Video Analysis

Use these tools to manage transcripts, retrieve chapter lists, run sentiment analysis, or start a new transcription job directly through your agent.

#	TOOL	DESCRIPTION
01	`delete_transcript`	Permanently removes a specific transcript record from your directory.
02	`get_chapters`	Retrieves an automated chapter list for media content.
03	`get_sentiments`	Analyzes the emotional tone of a transcript, identifying positive or negative moments.
04	`get_speakers`	Retrieves detailed labels separating and tracking different speakers in a conversation.
05	`get_summary`	Generates an automatic, concise summary of the full transcript content.
06	`get_topics`	Detects and lists the specific themes or topics discussed throughout the audio recording.
07	`get_transcript`	Checks the status of a transcription job or retrieves the completed transcript result.
08	`list_transcripts`	Shows you a list of your most recent and available transcripts for review.
09	`transcribe_audio_url`	Starts the process of transcribing any provided audio link.

See It in Action

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

U Transcribe this podcast URL: 'https://example.com/audio.mp3' and identify who speaks when.

Transcription Job Started.

The MCP has successfully queued the job. The high-fidelity transcript is being generated, including speaker labels.

Job ID: `tr_456`
Status: Processing (ETA: 2 minutes)

I've set up a monitor on this job and will notify you as soon as it's ready for review.

U Show me my five most recent jobs, and get the summary for 'Client Onboarding Call'.

Recent Transcripts Found:
- Team Sync (Completed)
- Product Interview (Completed)
- Marketing Strategy (Processing...)

I see the 'Client Onboarding Call' transcript. I'm running `get_summary` now...

Summary: The client confirmed interest in Feature X but needs pricing details for Enterprise Tier Y. They expressed concern about implementation timeframes, which we need to address in the next call.

U What was the general mood of the conversation on this audio? And what were the main topics covered?

Audio Intelligence Report (ID: tr_456):

- **Overall Sentiment:** Moderately Positive. Several segments showed high neutrality, especially during technical discussions.
- **Top Topics Detected:** 1. Implementation Roadmaps; 2. Pricing Models; 3. Feature X Adoption.

I've attached a segment breakdown showing where the sentiment dipped—it happened when discussing integration costs.

Frequently Asked Questions

01 **How do I use AssemblyAI MCP to transcribe video files from YouTube or Vimeo?**

You simply pass the public URL of the video to your AI agent. The MCP handles the streaming and transcription process, returning a full transcript that you can then analyze for summaries or topics.

02 **Can AssemblyAI MCP tell me who said what in an interview recording?**

Yes, it uses speaker diarization to label every utterance. You get detailed segments showing exactly which person spoke when, making meeting minutes accurate and easy to write up.

03 **What if I need the transcript for multiple recordings? Is there a way to process them all?**

Your agent can use the `list_transcripts` tool to see everything you've processed. From there, you can run analysis tools like getting summaries or topics on several jobs in sequence.

04 **Is AssemblyAI MCP better than just using a simple text-to-speech service?**

Yes, because it doesn't just transcribe; it analyzes the content. It pulls out insights like sentiment and topics, giving you deep context that basic transcription services miss completely.

05 **Can I use AssemblyAI MCP to organize my media library with chapters?**

Absolutely. The tool can automatically detect natural breaks in the audio or video and generate chapter markers (`get_chapters`), so you never lose your place when reviewing long content.

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 `"assemblyai-alternative": { "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

AssemblyAI 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 AssemblyAI. 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	AssemblyAI MCP
Server ID	`019dd0bd-7422-70bf-9913-8537412614bd`
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/assemblyai-alternative.