

MCP SERVER

NO CODE

CLOUD HOSTED

AssemblyAI MCP for AI Agents

Process and Audit Spoken Word from Audio Files

AssemblyAI provides a complete audio intelligence workflow for AI agents. It transcribes spoken content from any URL, giving you structured text that includes speaker labels and confidence scores. Manage job history, audit transcripts by sentence or paragraph, and ensure your audio data is always searchable and ready for analysis.

A+ Quality Score 98.33/100

speech-to-text

transcription

audio-processing

natural-language-processing

ai-models



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 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot 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 — connect your AI agent in under 60 seconds.

AssemblyAI MCP

6 tools available

Cloud-hosted on Vinkius

Connecting AssemblyAI to your agent transforms complex audio processing into a natural conversation. Instead of manually uploading files and waiting on web consoles, your agent handles the entire transcription process automatically. It starts jobs from any URL, retrieves clean text with speaker separation, and provides detailed audits on everything said. You can get transcripts broken down by sentences or paragraphs for structured data modeling, and even check confidence scores to verify accuracy. This level of audio intelligence management is available through Vinkius, the leading catalog of MCPs, allowing your agent to handle all media processing tasks without you ever needing technical access.

Whether you're monitoring a series of podcast episodes or transcribing lengthy meeting recordings, your agent acts as a real-time linguistic assistant. It monitors job status and maintains a full history of transcripts, keeping your audio assets organized and instantly searchable.

Core Capabilities

01 — Start Transcription Jobs

The MCP begins the process by taking an audio or video URL to initiate a new transcription job.

03 — Structure Text Data

The agent can break down the raw transcript into discrete paragraphs or individual sentences for precise data handling.

05 — Delete Records

The MCP allows you to delete specific transcript records when they are no longer needed.

02 — Retrieve Full Transcript Results

It fetches the complete written text, including speaker labels and confidence scores for every segment of speech.

04 — Monitor Job Status and History

You can list all past and active jobs, checking progress to ensure timely delivery of your audio content.

One Click on Vinkius — From Prompt to Execution

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

- 01** Connect your agent via any compatible client and provide the necessary API key.
- 02** Tell your agent which audio or video URL needs processing. The MCP starts a transcription job and monitors its progress.
- 03** Once complete, your agent retrieves the full text data, allowing you to structure it by paragraphs or sentences for immediate use.

The bottom line is that your agent manages the entire lifecycle: from starting the audio capture to delivering structured, verified text data without manual intervention.

Built For

Content creators and operations leads need this. If you spend time manually transcribing podcasts, auditing meeting minutes, or trying to turn recorded interviews into searchable databases, this MCP is for you. It takes the headache out of media data management.

Content Creator

You use it to monitor podcast transcriptions and instantly retrieve speaker metadata, letting your agent manage all published audio content.

Data Analyst

You connect it to verify transcription accuracy across multiple files and audit linguistic trends before feeding the data into a BI dashboard.

Operations Lead

You perform rapid audits of meeting records, asking your agent to pull out key summaries or specific actions items without reviewing the full transcript.

What Changes When You Connect

-
- 01 Get structured text output instantly. Use the `get_transcript_paragraphs` tool to break down long scripts into manageable blocks, perfect for database entry.

 - 02 Verify data quality immediately. The results include confidence scores, letting you audit every piece of transcribed speech and flag potential errors before publishing.

 - 03 Manage your media workflow without clicks. Your agent handles starting jobs (`transcribe_audio`), monitoring progress, and retrieving the final text—all through natural conversation.

 - 04 Maintain strict control over assets. Use `list_transcripts` to view a clean history of all job IDs and use `delete_transcript` when cleanup is necessary.

 - 05 Extract contextually rich data. By using `get_transcript_sentences` , you can build specific queries, asking your agent about details contained within precise conversational turns.
-

Real-World Applications

Analyzing Podcast Content for Show Notes

A content manager needs to turn a 45-minute podcast episode into searchable show notes. The agent uses `transcribe_audio` on the URL, then calls `get_transcript_paragraphs` and retrieves speaker labels to draft detailed, structured captions immediately.

Auditing Legal Meeting Minutes

An operations analyst receives a raw audio recording of a board meeting. They ask their agent to process it and then use the confidence scores from `get_transcript` to quickly flag any segments where transcription accuracy was questionable.

Indexing Academic Lectures

A data scientist wants to index a series of lectures for later retrieval. They run multiple jobs, using ``list_transcripts`` to manage the batch, and then use ``get_transcript_sentences`` to feed the structured data into an external knowledge base.

Cleaning Up Old Archives

A user realizes they have unnecessary old recordings. They ask their agent to run a cleanup command that calls ``list_transcripts`` and then systematically uses ``delete_transcript`` on expired jobs, keeping the system tidy.

Patterns to Avoid

Treating audio as simple text files

X AVOID

Copy-pasting a raw transcript into a spreadsheet and trying to manually separate speaker names or find specific timestamps. This is slow, error-prone, and ignores confidence levels.

✓ INSTEAD

Let your agent handle it. Use ``transcribe_audio`` first, then use the structured output from ``get_transcript`` which includes both speaker labels and confidence scores for precise data handling.

Ignoring job status

X AVOID

Assuming a long-running transcription is done because the initial API call succeeded. You might waste time processing incomplete or failed jobs.

✓ INSTEAD

Always check the progress using ``list_transcripts``. Wait until the agent confirms completion before attempting to retrieve data with ``get_transcript``.

Handling text in one piece

X AVOID

Receiving a massive wall of text and having no way to pinpoint the exact sentence where a key decision was made. You risk missing critical context.

✓ INSTEAD

Use ``get_transcript_sentences`` or ``get_transcript_paragraphs``. This breaks the content into small, targeted chunks that your agent can query for specific information.

The Right Fit

Use this MCP if your core workflow involves turning spoken word from audio files—podcasts, meetings, interviews—into structured, searchable text data. You need to audit accuracy (confidence scores) and break the content down by speaker or segment. Don't use it if you only need basic file conversion; other tools might handle simple WAV-to-TXT tasks. Also, don't use it if your primary goal is transcription from a live stream; this MCP requires pre-

recorded files via URL. If you just need to check metadata on existing records, `list_transcripts` helps, but for full content intelligence, the whole flow is necessary.

AssemblyAI MCP: Auditing Audio Content Accuracy

Today, turning audio recordings into usable text involves a painful cycle of manual listening, transcription service uploads, and then tedious copy-pasting. You wait for the file to process, download the resulting dump, and then spend hours cross-referencing speaker names or flagging sections where the software was unsure what it heard.

With this MCP, you simply ask your agent to run a job from a URL. It handles the whole pipeline, delivering not just the text, but also confidence scores for every segment. The result is clean, verifiable data ready for immediate use.

AssemblyAI MCP: Managing Spoken Word Data Structure

Before this, if you wanted to know what was said in a specific paragraph or by a single person, you had to rely on basic file search functions that often failed. You were limited to viewing the whole transcript as one giant block of text.

Now, your agent can use tools like `get_transcript_paragraphs` or `get_transcript_sentences`. This gives you granular control over the data structure—you get exactly what you need, broken down and ready for application logic.

AssemblyAI: 6 Tools for Audio Transcription and Auditing

These tools let your agent start jobs, retrieve structured text by sentence or paragraph, check job status, and manage transcript records.

#	TOOL	DESCRIPTION
01	<code>delete_transcript</code>	Removes a specified transcription record from the system's history.
02	<code>get_transcript_paragraphs</code>	Retrieves the full transcript text broken down into logical paragraphs.
03	<code>get_transcript_sentences</code>	Gets the transcribed content segmented and formatted by individual sentences.
04	<code>get_transcript</code>	Retrieves the final, processed text result of a completed transcription job.
05	<code>list_transcripts</code>	Lists all past and currently active transcription jobs in your account history.
06	<code>transcribe_audio</code>	Starts a new transcription job using any provided audio or video URL.

See It in Action

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

- U Can you transcribe this podcast episode from YouTube and give me key quotes with who said them?



Podcast Transcript Audit

Job ID: transcript-923

Status: Complete (Confidence Score: 98%)

Key Insights:

- Speaker A:** "The biggest shift this year is multimodal AI." (Confidence: 95%)
- Speaker B:** "It's not about the model; it's about the prompt engineering."
- Speaker A:** "Exactly. That's where we need to focus our resources next quarter." (Confidence: 97%)

- U List my last five transcript jobs and tell me which ones I need to clean up.



Transcript History Report

JOB ID	SOURCE TYPE	STATUS	DATE COMPLETED
t-102	Meeting Video	Complete	2024-06-18
t-923	Podcast Audio	Complete	2024-06-17
t-55a	Webinar Recording	Failed	N/A

The 'Webinar Recording' job is marked as failed. I recommend deleting it using the delete transcript tool to keep your list clean.

- U** I need the transcript from that meeting broken down by paragraphs, please. I want to analyze the flow.



Paragraph Breakdown for Meeting ID 501

The project scope needs immediate review.

The initial requirements are solid, but we missed a critical detail regarding international compliance.

Next steps involve drafting three separate proposals. We need one focused on the EU market and two others covering APAC regions. This requires more detailed input from legal.

Frequently Asked Questions

01 How do I find my AssemblyAI API Key?

Log in to your [**AssemblyAI dashboard**](https://www.assemblyai.com/app), and you will find your API Key on the main home page. Copy and paste it below.

02 What audio formats are supported?

AssemblyAI supports most common audio and video formats, including MP3, WAV, AAC, MP4, and others. Simply provide a public URL to the file.

03 Can the agent identify different speakers?







Yes. When starting a job via `transcribe_audio`, set the `speaker_labels` parameter to true. Your agent will return the text categorized by speaker ID.

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 <code>"assemblyai": { "url": "..."} </code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → 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](https://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	019d8418-9dd8-7398-a3a7-355fc0f5f6f5
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.