MCP SERVER | NO CODE | CLOUD HOSTED

iFLYTEK Open Platform MCP

Process voice and language data into structured insights.

iFLYTEK Open Platform provides deep speech intelligence right inside your workflow. Transcribe audio files to accurate text, generate synthetic voices from simple prompts, and analyze any body of text for keywords, sentiment, or key entities. It's a single source for advanced language processing—from real-time transcription to full document summarization.

A+ Quality Score 100/100

speech-to-text | voice-synthesis | nlp | audio-transcription | language-translation | linguistic-analysis

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/iflytek-open-platform

SCAN TO VIEW

The infrastructure that powers AI agents in the real world.

Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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.

iFLYTEK Open Platform / 讯飞开放平台 MCP

8 tools available
Cloud-hosted on Vinkius

This MCP lets your agent handle virtually any voice or natural language task without you having to jump through multiple web consoles. You can feed it audio files and have it accurately transcribe every word spoken. From that text, it can instantly summarize long documents, identify specific people or places using entity recognition, or even tell you if the original speaker sounded frustrated by analyzing sentiment.

If you need to generate content, simply give it a text prompt and get high-quality speech audio back. Need to talk to someone who speaks another language? You can translate entire conversations on the fly. By connecting iFLYTEK through Vinkius, your agent becomes a real-time linguistic assistant, acting as a unified layer that handles everything from raw audio capture to structured data extraction, all within your preferred AI client.

Core Capabilities

01 — Transcribing Audio

It converts spoken words from any audio file or live stream into clean, usable text.

02 — Analyzing Text Meaning

Your agent can extract key terms, identify people and places, and determine the overall emotion (sentiment) of a piece of writing.

03 — Translating Languages

It translates text between multiple languages instantly, making international communication simple for your agent.

04 — Synthesizing Speech

You can generate high-quality speech audio from any written text using customizable voice models.

05 — Extracting Information

It handles complex tasks like summarizing large documents or pulling readable text directly out of images (OCR).

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/iflytek-open-platform** — connect your AI agent in three steps.

01 Subscribe to this MCP and provide your iFLYTEK App ID, API Key, and Secret.

02 Give your AI client a prompt that references the audio or text data you need processed.

03 Your agent calls the necessary tool (like `speech_to_text` or `summary_generation`) and returns structured, actionable results.

The bottom line is: instead of using multiple specialized web tools, your AI client runs all these language functions through one single endpoint.

Built For

Content producers who deal with multilingual media; customer support analysts drowning in call transcripts; and developers building advanced multi-stage agent pipelines.

Customer Support Analyst

They feed the MCP raw audio recordings of customer calls, then use `text_sentiment` to flag negative interactions and `summary_generation` to create tickets for follow-up.

Content Creator / Editor

They upload a video file, run `speech_to_text` on it, then use the resulting text to generate captions, translate it into three languages using `translate`, and finally create short social media snippets.

Developer / Engineer

They integrate this MCP into a larger agent framework, allowing their code to process everything from image-based receipts (`ocr_general`) to transcribed conversations in a single workflow.

What Changes When You Connect

01 You gain immediate insight from audio. Instead of manually transcribing a meeting recording, simply use `speech_to_text` to get clean text that your agent can immediately analyze for key findings.

02 Automate content repurposing. Use the MCP to take a long-form article (via `summary_generation`), translate it (`translate`), and then convert the resulting summary into an audio file using `text_to_speech`—all in one go.

03 Improve customer handling by analyzing tone. Feed call transcripts into the tool, and `text_sentiment` instantly flags any conversation that shows high levels of negative emotion, letting you prioritize urgent follow-up.

04 Read text from anything. If your input includes receipts or handwritten notes, use `ocr_general`. This bypasses manual data entry entirely by turning images into usable, structured text for the agent to process.

05 Keep language barriers out of the loop. When dealing with global teams, running `translate` ensures that everyone gets consistent, accurate meaning without requiring human interpretation at every stage.

Real-World Applications

Analyzing International Customer Calls

A support manager records 50 multilingual calls daily. Instead of manually reviewing transcripts, they prompt their agent to use `speech_to_text` and `translate` on all files first. Then, running `entity_recognition` isolates the customer's account ID and geographic location from every single interaction.

Creating Multilingual Video Content

A marketing team records a core message. They use the MCP to run `speech_to_text` on the master recording, then pass that text through `summary_generation` for short clips. Finally, they repeat this process using `translate` and `text_to_speech` for three different market voices.

Auditing Legal Documents

A compliance officer receives dozens of PDFs with handwritten notes or stamps. They run `ocr_general` to extract all visible text into the agent. Then, they use `keyword_extraction` and `text_sentiment` on the resulting data to check for specific risk indicators.

Building an Agent Dashboard

A developer wants a single dashboard that accepts audio files. The agent uses `speech_to_text` first, then passes the text through `text_sentiment`. This gives the dashboard not just the transcript, but also a real-time 'Risk Score' based on language tone.

Patterns to Avoid

Treating it like simple search

✗ AVOID

Trying to use `text_sentiment` just because you want to know if a document is about tech or finance. It won't categorize topics; it only measures tone.

✓ INSTEAD

If your goal is classification, run `keyword_extraction` first. If the keywords are consistent (e.g., 'chip,' 'GPU,' 'AI'), then you can determine the topic.

Skipping audio preprocessing

✗ AVOID

Giving the MCP a raw, noisy recording and expecting perfect results without checking for language or file format compatibility.

✓ INSTEAD

Always run `speech_to_text` first. If you know the language, specify it in the prompt to maximize transcription accuracy.

Over-relying on one tool

✗ AVOID

Using only `summary_generation` and missing critical details because you didn't extract specific names or dates.

✓ INSTEAD

Always pair summarization with `entity_recognition`. This ensures the summary is accurate *and* retains actionable data points like names, places, and organizations.

The Right Fit

Use this MCP if your primary bottleneck involves turning unstructured language—be it spoken audio, scanned images, or multi-lingual text blocks—into structured, analyzed insight. If you need to perform translation alongside transcription, or analyze the emotional tone of a document, this is your tool. However, don't use it if your core task is simply data storage or retrieval from a known database (use a dedicated database connector instead). Also, remember that while `keyword_extraction` helps identify topics, you

must run it *before* relying on those keywords for complex decisions; otherwise, the results are just educated guesses.

Manually processing global communications is a nightmare.

Today's workflow means logging into separate services: one for transcription, another for translation, and yet a third to analyze customer mood. You record an audio file, download the transcript, copy it into a translator, then paste that result into a sentiment analysis dashboard. This cycle of copying, pasting, renaming files, and waiting is slow and introduces human error at every single handoff.

With this MCP, you give your agent the raw audio file once. It handles the transcription (`speech_to_text`), translates it to English (`translate`), and immediately runs `text_sentiment`—all in one call. You get a clean JSON object containing the transcript, the translated text, and the sentiment score, ready for immediate use.

Getting actionable insights using entity_recognition.

Without this tool, you read through 30 pages of meeting notes. You're looking for specific names, project codes, and dates that need to be logged into a CRM. You highlight them manually, copy the text, and paste it into a spreadsheet, praying you didn't miss anything.

Now, your agent processes the document through `entity_recognition`. It doesn't just give you text; it gives you structured data—a list of recognized names, locations, and dates—that feeds directly into your database. You move from reading to acting instantly.

iFLYTEK Open Platform: 8 Tools Available

These tools allow you to process spoken word data, extract key information from documents, translate languages, and generate synthetic audio content through your AI agent.

#	TOOL	DESCRIPTION
01	entity_recognition	This tool identifies and extracts specific names, locations, or other defined entities from any body of text.
02	keyword_extraction	It pulls out the most important topics or core vocabulary words that describe a given piece of text.
03	ocr_general	This tool reads and extracts digital text from images, allowing you to process documents that aren't typed out.
04	speech_to_text	It transcribes spoken words recorded in audio files into accurate written text.
05	summary_generation	You feed it a large block of writing, and it returns a concise summary hitting all the main points.
06	text_sentiment	This tool analyzes written text to determine if the tone is positive, negative, or neutral.
07	text_to_speech	It converts any given block of written text into high-quality synthetic speech audio.
08	translate	This tool translates written text accurately between many different world languages.

See It in Action

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

U Convert this text to speech: 'Hello, welcome to the world of AI!'.

I've synthesized the voice for you. The audio data has been generated and is ready for use. Would you like to use a specific voice model or change the speech rate?

U Translate 'Artificial Intelligence' from English to Chinese.

I've translated your text. 'Artificial Intelligence' in Chinese is '人工智能' (Rén gōng zhì néng). Should I also analyze the sentiment of this phrase or extract keywords?

U Summarize this long text about future tech trends: [Text].

I've generated a summary. The main points are: 1) AI integration in daily hardware, 2) Growth of edge computing, and 3) Decentralized cloud storage. Would you like me to identify the key entities mentioned in the full text?

Frequently Asked Questions

01 **How does iFLYTEK Open Platform handle different languages?**

It handles multiple languages through dedicated tools like `translate` and the core transcription engine. You simply specify the source language in your prompt, and it manages the complexity.

02 **Can I use iFLYTEK Open Platform to read text from photos?**

Yes, you can run `ocr_general`. This tool reads images—like signs or handwritten notes—and converts them into plain, machine-readable text that the agent can then process.

03 **What is the difference between keyword_extraction and entity_recognition?**

`keyword_extraction` pulls out general topics (e.g., 'AI,' 'market trends'). `entity_recognition` finds specific, named things, like a person's name ('John Smith') or a company ('Google LLC').

04 **Is the text_to_speech audio high quality?**

The generated speech is customizable and designed to be high quality. You can specify different voice models or adjust parameters in your prompt.

05 **Does iFLYTEK Open Platform require me to manually summarize the text?**

No, you use the `summary_generation` tool. Your agent handles the summarization process based on how much detail you ask it to retain in the prompt.

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 `"iflytek-open-platform": { "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

iFLYTEK Open Platform / 讯飞开放平台 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 iFLYTEK Open Platform / 讯飞开放平台. 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	iFLYTEK Open Platform / 讯飞开放平台 MCP
Server ID	`019d8447-e18e-716b-a825-90ecd8c1fa72`
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/iflytek-open-platform.