MCP SERVER | NO CODE | CLOUD HOSTED

FileStack MCP

Analyze, Transform, and Process Media Files.

FileStack lets your AI agent handle complex media workflows end-to-end. Upload files from public URLs, extract text using OCR, analyze images for tags and safety status, generate optimized CDN links, or start video transcoding jobs—all through simple commands.

A+ Quality Score 100/100

file-upload | image-processing | ocr | video-transcoding | content-analysis

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026

vinkius.com/mcp/filestack

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.

FileStack MCP

8 tools available
Cloud-hosted on Vinkius

This MCP connects advanced file handling capabilities directly into your workflow. You can automate asset pipelines that previously required multiple manual steps: uploading media from a public web link, extracting structured text from scanned documents, and analyzing images for visual content tags. Need to turn a raw video file into web-ready formats like MP4? You can initiate those complex transcoding jobs without ever touching a command line. Furthermore, you never have to manually resize or crop an image again; the system generates optimized URLs on demand. If your current AI agent setup is running into bottlenecks managing diverse media types, connecting this MCP via Vinkius gives you immediate access to these powerful tools—everything from basic file inspection to advanced video processing.

Core Capabilities

01 — Analyze image content

Detects objects and features within an uploaded image.

02 — Extract text from media

Pulls printed or handwritten text out of documents and images using OCR technology.

03 — Check content safety

Determines if an image contains unsafe or inappropriate material.

04 — Generate optimized media links

Creates transformation URLs that allow resizing, blurring, or filtering of images without manual edits.

05 — Manage video transcoding jobs

Starts and monitors background processes to convert videos into various web-ready formats.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP, providing your unique Filestack API Key (and optional policy details for extra security).

02 Your AI client uses the key to communicate with the service endpoint and initiate a file operation.

03 The system returns status updates or the final assets—whether that's an extracted text block, a video job UUID, or a transformation URL.

The bottom line is you tell your agent what kind of media task to run, and it handles the entire technical process.

Built For

Anyone dealing with high volumes of visual assets or varied file formats needs this. Think content moderators who need scale, developers building complex frontends, or data engineers needing structured text from unstructured images.

Content Moderator

Uses the MCP to automatically scan and flag thousands of visual assets for unsafe content or required tagging before they go live.

Full-Stack Developer

Automates asset pipelines by generating transformation URLs or uploading images directly from a public web link within their code editor.

Data Engineer

Extracts structured text and metadata from scanned documents, allowing the data to feed into other analytical systems.

What Changes When You Connect

01 You skip manual image editing. Instead of manually resizing files for different platforms, calling `generate_transform_url` instantly generates the correct URL with specific dimensions or filters applied.

02 Content safety becomes automatic. Before publishing anything, your agent uses `get_sfw_status` to check for unsafe content, giving you immediate compliance feedback at scale.

03 Complex video formats are handled in the background. Initiate and monitor entire transcoding pipelines using `start_video_transcode`, converting one source file into multiple web-ready versions like MP4 or HLS.

04 Data extraction is simple. If you receive a scanned document, running `get_ocr` pulls the raw text instantly, letting your agent work with structured data without manual input.

05 Speed up uploads. Don't worry about local file paths; using `upload_from_url` lets your agent pull and store assets directly from any public web link.

Real-World Applications

Processing a batch of user-submitted photos

A content moderation team receives 500 new profile pictures. Instead of reviewing them one by one, the agent runs `get_sfw_status` and `get_image_tags` on every file. The results are compiled into a report showing which files need human review versus those that pass automatically.

Building a dynamic article landing page

A developer needs an image displayed at 400px wide and blurred for preview purposes. They use `generate_transform_url` to get the exact link, ensuring the front-end only loads the correct, optimized version of the asset.

Digitizing old archival documents

A data scientist has a collection of scanned annual reports. The agent uses `get_ocr` on each file to pull out all text, then runs `get_metadata` to understand the original dimensions and source format before feeding the clean text into a database.

Preparing a YouTube video for multiple platforms

A marketing manager uploads a high-res master video. The agent calls `start_video_transcode` to create versions optimized for Instagram Reels, web embeds, and mobile viewing. They then use `get_video_status` until all jobs are complete.

Patterns to Avoid

Trying to manually resize images

✗ AVOID

When building a site that needs 10 different image sizes, the developer has to write code for each size and maintain complex file paths.

✓ INSTEAD

Use `generate_transform_url`. This single tool lets your agent generate an optimized URL for any specific size or filter without you ever having to touch the source asset again.

Assuming all images are clean text

✗ AVOID

A user tries to analyze a scanned textbook page using a simple OCR library that only reads digital PDFs, failing when confronted with handwritten notes.

✓ INSTEAD

Run `get_ocr`. This tool is specifically designed to handle the complexity of extracting readable text from both printed and handwritten materials within images.

Overlooking file details

✗ AVOID

A developer uploads a file but forgets if it's actually a JPEG, or what its original dimensions were, leading to layout errors.

✓ INSTEAD

Call `get_metadata`. This tool retrieves all the underlying technical specifics about any asset—dimensions, mime type, and size—giving you complete context upfront.

The Right Fit

Use this MCP if your workflow involves managing diverse media types, specifically images, videos, or scanned documents. If you need to analyze visual content for tags or safety, generate optimized links, or process text out of pictures, this is the tool. Don't use it if you only need basic cloud storage functionality; other general-purpose file tools handle simple uploads. Also, don't use it just because your AI client can talk to a bunch of APIs—you must have a genuine media processing task (like OCR or transcoding) that requires this specific level of deep analysis.

Dealing with media assets is usually a mess of manual conversions and checks.

Right now, every time you upload an image for use online, you probably have to check its size, crop it for Instagram, then maybe resize it again for your website header. If that asset was video, you're likely juggling three separate tools just to get a usable web link, and if it came from a scan, you're copy-pasting text block by manual text block.

With this MCP connected through Vinkius, the process becomes one command. Your agent handles all those checks—size, safety, format—and gives you back exactly what you need, whether that's an optimized URL or clean, extracted text.

FileStack enables comprehensive asset management with its tools.

The separate steps of checking file metadata, running object detection via `get_image_tags`, and then initiating a video conversion are all separated by manual handoffs. You check the size in one system, upload to another, and run the processing job somewhere else.

Now, your agent handles it all sequentially. It checks the metadata with `get_metadata` first, uses that information to determine if `generate_transform_url` is needed, and then executes the transformation—all within a single conversation flow.

FileStack: Media Processing Tools (8 tools)

These tools let you perform specific media operations like checking metadata, extracting text, generating transformation links, and managing video conversion jobs.

#	TOOL	DESCRIPTION
01	`get_image_tags`	Analyzes an image and returns a list of detected objects or features.
02	`get_metadata`	Pulls deep technical details about a file, including its dimensions, size, and mime type.
03	`get_ocr`	Scans an image or document to pull out any visible printed or handwritten text.
04	`get_sfw_status`	Checks whether uploaded content violates safety standards and is safe for public viewing.
05	`generate_transform_url`	Builds a URL that, when accessed, will automatically resize or filter an image to specific parameters.
06	`upload_from_url`	Copies and uploads any file found at a public web address into your managed storage.
07	`get_video_status`	Checks the current progress of an ongoing video transcoding job using its unique ID.
08	`start_video_transcode`	Initiates a background process to convert an audio or video file into different web formats.

See It in Action

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

U Upload this image to Filestack: https://example.com/photo.jpg

I've uploaded the file. The new Filestack handle is `ABC123XYZ`. You can now use this handle for transformations or metadata analysis.

U What objects are detected in the image with handle ABC123XYZ?

Analyzing the image... I've detected several tags: 'mountain' (98% confidence), 'snow' (95%), and 'outdoor' (92%).

U Convert the video ABC123XYZ to mp4 format.

I've started the transcoding job for handle ABC123XYZ using the 'mp4' preset. The job UUID is `job-550e8400`. You can check the status later using `get_video_status`.

Frequently Asked Questions

01 **How do I process text from an image using FileStack?**

You use the `get_ocr` tool. Simply pass the image handle to your agent, and it will extract all visible printed or handwritten text into a clean, usable string for you.

02 **Can I automatically resize an image using FileStack?**

Yes, use `generate_transform_url`. This tool does not change the file; it just gives you the specific URL needed to serve that image at a precise width and height.

03 **What is the difference between uploading files and using upload_from_url?**

Uploading manually handles local files, but `upload_from_url` lets your agent grab an asset directly from any public web link, simplifying workflows that rely on external content.

04 **Does FileStack handle video formatting for me?**

Yes. You use `start_video_transcode` to tell the system what format you need (like MP4). The agent then monitors the progress using `get_video_status` until it's done.

05 **Does FileStack check if my content is safe for work?**

It does. You run the `get_sfw_status` tool, and the agent will immediately tell you whether the image passes or fails safety checks before publication.

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 `"filestack": { "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

FileStack 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 FileStack. 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	FileStack MCP
Server ID	`019e3896-e3fb-7306-ae38-cd0230e64c2b`
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/filestack.