MCP SERVER | NO CODE | CLOUD HOSTED

MonkeyLearn MCP

Stop guessing what customers mean from raw text.

MonkeyLearn analyzes text data for deep insights using natural language processing via your AI client. It classifies sentiment, pulls out keywords, detects topics from customer reviews or articles, and extracts specific information directly into a structured format.

A+ Quality Score 100/100

natural-language-processing | sentiment-analysis | text-classification | data-extraction | machine-learning

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

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.

MonkeyLearn MCP

10 tools available
Cloud-hosted on Vinkius

Stop treating raw text like an unorganized pile of notes. This MCP lets you analyze messy text—customer feedback, article content, support tickets—and pull out the actual data points that matter using your AI client. You don't need to write custom Python code or build complex ETL pipelines just to understand what people are saying.

It handles everything from basic sentiment checks (is this positive or negative?) to identifying specific entities and topics you haven't even thought of yet. If you can describe the data, your agent can find it. You connect this MCP through Vinkius, giving your AI client access to a full library of text analysis tools alongside anything else you use in the catalog.

Core Capabilities

01 — Determine Text Sentiment

It classifies whether text is positive, negative, or neutral, giving you a confidence score for each rating.

02 — Extract Specific Data Points

The agent pulls out named entities, keywords, and structured data (like product names or dates) from unstructured blocks of text.

03 — Categorize Content Type

You can assign content to predefined topics or intent categories using pre-trained models.

04 — Discover Available Models

The MCP allows you to list and inspect all the specialized classification and extraction tools available in your account.

05 — Review Workflow History

It tracks automated processing activity, letting you see how many times text has been analyzed over a given period.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and paste your unique MonkeyLearn API Key into the Vinkius connection settings.

02 Use your AI client to provide the raw text you want analyzed, along with instructions on what kind of data you need (e.g., 'Find all product names and the sentiment').

03 Your agent sends the request through this MCP, which returns a clean, structured JSON output containing the extracted keywords, topics, or sentiment scores.

The bottom line is that it takes unstructured text input from your AI client and gives you organized data outputs like spreadsheets.

Built For

This MCP is for anyone who spends time reading customer feedback, support tickets, or market research reports but doesn't want to hire a full-time NLP team. It targets Product Managers tired of manual spreadsheet analysis and Content Teams needing automated SEO data.

Product Manager

They feed thousands of customer reviews into the agent, which uses classification tools to identify the top three most complained-about features or topics.

Content Strategist

They run articles through the MCP to automatically pull out key phrases and tag structures for SEO planning without manual keyword research.

Data Scientist

They quickly test custom machine learning models against sample data streams, getting results instantly without writing boilerplate code.

What Changes When You Connect

01 Identify immediate pain points. Instead of manually reading thousands of reviews, you use classification tools to automatically flag high volumes of negative sentiment or specific topics using `classify_text`.

02 Structure messy data instantly. The agent doesn't just summarize; it runs the `extract_text` tool to pull out structured lists—like every unique product code mentioned in a week's worth of support tickets.

03 Manage your models via chat. You don't need to navigate complex web dashboards. Simply ask the MCP to list all available tools using `list_classifiers` or `list_extractors` and start building.

Real-World Applications

Analyzing a Product Launch Wave of Reviews

A PM needs to know if customers like the new UI. Instead of reading 500 reviews, they ask their agent to process them using `classify_text`. The MCP returns a breakdown: 60% positive sentiment, with key topics like 'Navigation' and 'Speed' being mentioned most often.

Monitoring Support Ticket Trends

A support lead needs to know if a specific bug is spiking. They feed the last month of tickets into their agent, which uses `list_classifiers` to run topic detection and alerts them immediately when 'Login failure' exceeds 15% of all incoming text.

Competitive SEO Keyword Harvesting

A content team needs keywords for a new article. They provide the competitor's page URL text, and the agent uses `extract_text` to pull out every specialized term or entity mentioned, saving hours of manual research.

Debugging Automated Processes

A data analyst wants to know why a recent workflow failed. They use the MCP's ability to list workflows and check `list_activity` to pinpoint exactly which step or model caused the processing failure in the last hour.

Patterns to Avoid

Using simple search/find features.

✗ AVOID

Trying to manually copy 10,000 customer reviews into a spreadsheet and using conditional formatting to color-code 'negative' entries. You miss nuance, and it takes hours of tedious work.

✓ INSTEAD

You connect the MCP and ask your agent to run `classify_text` against the entire batch. The result is an immediate, structured column with sentiment scores for every single review.

Over-relying on general summarization tools.

✗ AVOID

Asking a generic AI to 'summarize these articles.' It gives you fluffy paragraphs but fails if you need specific data points like dollar amounts, names, or dates mentioned across the whole set.

✓ INSTEAD

Use `extract_text` instead. This tool forces the agent to find and output only specific, named entities from the text into a clean list.

Building custom code for every project type.

✗ AVOID

Writing unique Python functions just to check sentiment or pull out keywords for one single client. This means maintaining 10 different pieces of brittle codebases.

✓ INSTEAD

Connect this MCP through Vinkius. Your agent handles the complexity, letting you access proven tools like `list_classifiers` and applying them via natural language conversation.

The Right Fit

Use this MCP if your core problem is understanding meaning, context, or structure within massive amounts of unstructured text. Specifically, use it when you need to move beyond simple keyword counts and actually understand the *intent* behind the words (sentiment analysis) or pull out specific facts that were never highlighted (data extraction). You should also use `list_classifiers` if you are unsure which type of analysis is needed. Don't use this MCP if your data problem involves images, audio files, or real-time sensor feeds; for those tasks, you need a different category tool. If all you need is to send an email or check a database status, use a dedicated messaging or CRM connector instead.

The pain of manually sifting through feedback streams.

Think about it: your team gets dozens of support tickets and hundreds of customer reviews daily. The current process is painful copy-pasting text into spreadsheets, then spending hours manually skimming every entry to find patterns—'Oh, look, three people mentioned 'login button' today.' It's slow, exhausting, and you inevitably miss the subtle trends.

With this MCP, your agent handles the grunt work. You just feed it a batch of text, tell it what kind of insight you need, and it spits out clean data lists—like a count of every topic found or a clear list of all names mentioned. The hard part is done automatically.

Getting structured insights with MonkeyLearn MCP

The manual steps that disappear are the filtering, the grouping, and the human interpretation. You no longer need to spend time building custom scripts just to check if a review is positive or negative. The agent does the classification using `classify_text` instantly.

What's different now is speed and scale. You move from spending days manually analyzing feedback to getting immediate, actionable data reports in minutes.

MonkeyLearn: 10 Text Analysis Tools

Use these ten tools to classify text, extract specific data points, check workflow status, and understand the full capabilities available in your MonkeyLearn account.

#	TOOL	DESCRIPTION
01	`classify_text`	Uses a specific model to assign a category or topic label to a piece of text.
02	`extract_text`	Pulls out structured data, like names, dates, and key phrases, from unstructured text.
03	`get_classifier_details`	Fetches detailed configuration information about a specific classification model by its ID.
04	`get_extractor_details`	Retrieves the metadata and setup details for an extraction tool using its unique identifier.
05	`list_activity`	Provides a summary of all text processing runs that have occurred in your account recently.
06	`list_classifiers`	Displays every available classification tool, such as sentiment analysis or topic detection, set up in your account.
07	`list_extractors`	Lists all available data extraction tools, including keyword and entity recognition modules.
08	`list_pipelines`	Shows a list of automated workflows that process text on your behalf.
09	`list_tag_trees`	Retrieves the full hierarchy and structure of tags used by your classification models.
10	`list_workflows`	Shows a list of currently running or configured automated data processing jobs.

See It in Action

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

U Classify the sentiment of this review: 'The product exceeded all my expectations, truly amazing!' using model cl_oZ9GRg8P.

I've analyzed the text. The classifier cl_oZ9GRg8P identified the sentiment as 'Positive' with a confidence score of 98.5%.

U List all classifiers available in my account.

Retrieving your classifiers... I found 4 models: 'Sentiment Analysis' (cl_oZ9GRg8P), 'Topic Detection' (cl_piuz8L9), 'Support Ticket Classifier', and 'Email Intent'. Would you like to check the details for any of these?

U Show me my recent processing activity.

I've fetched your account activity. In the last 24 hours, you've made 156 API calls across 3 different models, mostly using the 'Sentiment Analysis' classifier. You have plenty of credits remaining for the current period.

Frequently Asked Questions

01 **Where do I find my MonkeyLearn API Key?**

Log in to MonkeyLearn and go to your API Settings page. You'll find your personal API Key there.

02 **Can I use custom models I've trained myself?**

Yes! The `classify_text` and `extract_text` tools work with both pre-trained models and any custom models you have created in your account. Just provide the specific Model ID.

03 What is a Pipeline in MonkeyLearn?

Pipelines allow you to chain multiple processing steps together (e.g., classification followed by extraction). You can use the `list_pipelines` tool to see what's available in your account.

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 `"monkeylearn": { "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

MonkeyLearn 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 MonkeyLearn. 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	MonkeyLearn MCP
Server ID	`019d75d8-2698-7169-8423-df71589568df`
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/monkeylearn.

© 2026 Vinkius · vinkius.com · support@vinkius.com