

MCP SERVER

NO CODE

CLOUD HOSTED

Helpjuice MCP

Manage Documentation, Articles, and Analytics via Chat.

Helpjuice connects your knowledge base directly to any AI agent, letting you manage all documentation tasks through conversation. Use this MCP to list articles, create new content drafts, update existing guides, and run detailed analytics on usage trends without touching a CMS dashboard. It gives your AI client full control over internal and external support documentation.

A+ Quality Score 100/100

knowledge-base

documentation

self-service-support

article-management

search-optimization

internal-wiki



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.

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.

Helpjuice MCP

12 tools available

Cloud-hosted on Vinkius

You can treat your entire knowledge base like an extension of your chat window. Instead of clicking through complex folders or filling out forms in a separate system, you just talk to your agent about what you need done with your articles.

Want to write a new guide? Just ask the agent; it drafts and saves the article for you. Need to know if an article is popular? You can pull up engagement statistics instantly. If a user searches something that yields zero results, the AI can check recent search trends to tell you what content is actually missing.

The system handles everything from listing all your categories to checking which internal users have permission to edit guides. This makes your AI client act less like an assistant and more like a dedicated documentation specialist. Accessing this functionality through Vinkius means connecting one time gives your agent instant, deep access to the whole catalog of knowledge tools.

Core Capabilities

01 — Search Documentation

Find existing answers across your entire knowledge base using natural language queries.

02 — Create and Modify Articles

Draft brand new articles or make changes to the body and metadata of guides directly from chat prompts.

03 — Analyze Content Usage

Pull up key metrics, like view counts and user voting records, for specific pieces of documentation.

04 — Manage Structure

List and organize your content into defined categories, giving you a map of all available guides.

05 — Monitor Content Health

Review recent search terms or list contributors to identify gaps in documentation coverage.

One Click on Vinkius — From Prompt to Execution

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

- 01 Subscribe to this MCP and provide your Helpjuice API Key and Subdomain credentials.
- 02 Your AI client connects to the service, authorizing access to all knowledge base tools.
- 03 You give a command in chat—for example, 'Find articles about X'—and the agent uses the appropriate tool to fetch or modify data.

The bottom line is, you never have to navigate your documentation platform manually again; your AI client does it for you.

Built For

This MCP is essential for Knowledge Managers who are tired of manual reporting and Support Agents who need instant access to the right guide during a customer call. If your team spends too much time copying article content or manually checking engagement stats, you need this.

Knowledge Manager

They use this MCP to list articles and check categories, ensuring documentation is structured and that no guides are outdated.

Support Agent

They rely on the agent's search capability to quickly find troubleshooting answers during a live customer interaction.

Technical Writer

They use this MCP to create articles, update drafts, and track content helpfulness by reading upvotes or recording downvotes.

What Changes When You Connect

- 01 Stop digging through CMS menus. Instead of manually listing articles using `list_articles`, just ask your agent for 'all published guides' and get the list instantly.

- 02 Quickly spot content gaps by running `get_search_trends` . If users are searching for 'API limits' but no article exists, you know exactly what to write next.

 - 03 Content changes become simple. You can update an entire guide using `update_article` without ever needing to log into the backend editor.

 - 04 Track content performance with one command: `get_article_stats` . See which guides are viewed often and which ones need a refresh.

 - 05 Improve quality control by running user checks, like listing contributors via `list_kb_users` , so you know who owns what documentation section.
-

Real-World Applications

A guide is falling out of date and needs an overhaul.

The writer asks the agent to check `get_article_stats` for 'Billing FAQs'. The stats show low engagement and many downvotes. The writer then uses `get_article_details` to pull the old content, drafts a new version, and finally uses `update_article` to push the fix.

A new product feature needs to be documented and categorized.

The creator first uses `list_categories` to see where the guide fits. They then run `create_article` for the draft, specifying the title and body content before requesting the agent to list all articles again to confirm placement.

The Support team needs to know what documentation is missing.

Instead of having agents manually check user feedback against existing guides, the manager runs `get_search_trends` . The agent reports that 'OAuth error codes' is a top search term with no matching articles, allowing immediate content creation.

Checking permissions across different teams.

The ops engineer uses `list_kb_groups` to see who can edit guides. Then they use `list_kb_users` to check specific team members against those groups, ensuring proper access control before deployment.

Patterns to Avoid

Trying to find a guide using general chat.

X AVOID

Asking 'How do I fix my account?' and getting irrelevant search results because the AI can't narrow down the scope or check for specific missing articles. This wastes time and frustrates the user.

✓ INSTEAD

Use ``search_kb`` with precise terms, like 'account password reset guide'. If that fails, run ``get_search_trends`` to confirm if the topic is a common search gap.

Manually copying and pasting articles.

X AVOID

A content writer has to open three different tabs in the CMS just to gather text snippets for a single new guide, leading to copy/paste errors and wasted time.

✓ INSTEAD

Use ``get_article_details`` or ``list_articles`` to pull all necessary source material directly into your chat window. Then use ``create_article`` to draft the final piece.

Guessing which article needs updating.

X AVOID

A manager sees a high number of downvotes but doesn't know if it's because the content is wrong or just hard to read. They waste time guessing at the root cause.

✓ INSTEAD

Use ``get_article_stats`` first. This gives you specific numbers on views and votes, confirming both the visibility and the helpfulness rating.

The Right Fit

You need this MCP if your daily job revolves around keeping documentation accurate, searchable, and up-to-date. If you frequently have to manually check content gaps, pull usage data, or update guides across multiple backend systems, this is for you. Use it when you want your AI agent to act as a full Knowledge Manager, able to list articles, create drafts, and analyze stats using `get_article_stats`.

Don't use this if all you need is general chat support or simple document storage without structured metadata. If your primary goal is just file sharing or basic Q&A that doesn't require tracking authorship or content versions, a simpler search tool will suffice. You only need the depth of article management this MCP provides.

The documentation chaos that slows down support teams.

Right now, when a customer calls with a tricky question, your agent can't just magically pull up the perfect answer. They have to navigate three different tabs: one for search results, one for user groups to see who is allowed to edit, and another dashboard showing usage stats that might tell you if the guide even exists.

With this MCP, all of that complex navigation disappears. Your AI agent handles it behind the scenes. You ask a question about content or status, and you get a clean, actionable answer right in your chat window.

Helpjuice MCP: Control every piece of knowledge.

No more hunting through different CMS folders to check if an article is live or needs revisions. You can instantly list all articles, get the full content for `get_article_details` , and even track who last edited it using `list_kb_users` .

It's not just about finding things; it's about managing them completely. This MCP gives you conversational control over your entire knowledge base lifecycle.

Helpjuice: 12 Tools for Documentation Management

These tools allow you to perform every core function of knowledge base management—from drafting new guides to analyzing usage data—all within your AI agent.

#	TOOL	DESCRIPTION
01	<code>create_article</code>	Generates a brand new article title and body within the knowledge base.
02	<code>downvote_article</code>	Records that an article was unhelpful to the user.
03	<code>get_article_details</code>	Retrieves all metadata and the full text content for a specific guide.
04	<code>get_article_stats</code>	Pulls up key usage data, including views and votes, for one article.
05	<code>get_search_trends</code>	Lists the most common search terms used by site visitors over a period of time.
06	<code>list_articles</code>	Provides an overview list of every article in your knowledge base.
07	<code>list_categories</code>	Shows all the main organizational sections used to group articles.
08	<code>list_kb_groups</code>	Retrieves a list of internal user permission groups defined in the system.
09	<code>list_kb_users</code>	Lists all internal contributors and authors who have access to edit guides.
10	<code>search_kb</code>	Searches the entire knowledge base content for articles matching a given query.
11	<code>update_article</code>	Modifies the title or body of an existing article guide.
12	<code>upvote_article</code>	Records that an article was helpful to the user.

See It in Action

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

U Find articles related to 'OAuth integration'.



Searching knowledge base... I found 3 relevant articles: 'Setting up OAuth 2.0', 'API Authentication Guide', and 'Troubleshooting Connection Issues'. Would you like the full content for the 'Setting up OAuth 2.0' article?

U List all categories in my knowledge base.



I've retrieved your categories. You have 5 main sections: 'Getting Started', 'User Guides', 'Developer API', 'Release Notes', and 'FAQs'. Should I list the articles inside the 'Developer API' category?

U Show me the engagement stats for article ID 1021.



Retrieving analytics... Article ID 1021 ('Advanced Filtering') has had 450 views and 32 upvotes this month. It is currently one of your top-performing pieces of documentation. Would you like to see recent search terms that led users here?

Frequently Asked Questions

01 How does Helpjuice help with article creation?

You can use the `create_article` tool to generate a new draft guide directly from chat. This bypasses opening any content management system (CMS) and saves time.

02 Can I check if an article is popular using Helpjuice MCP?

Yes, use `get_article_stats` to pull up engagement statistics. You'll see detailed metrics like view counts and how many users gave it a helpful vote.

03 How do I find out what questions customers are asking that aren't covered?

Run the `get_search_trends` tool to list recent search terms. This shows you exact gaps in your knowledge base, guiding where you need to write new content.

04 Does Helpjuice MCP let me organize my guides into sections?

Yes, you can use `list_categories` to see your current structure and then manage the grouping of articles through the agent's conversational commands.

05 If I update a guide, do I need to manually publish it again with Helpjuice MCP?







The MCP handles the publishing flow. You simply ask to `update_article`, and the system manages pushing that revision live within your knowledge base.

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>"helpjuice": { "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

Helpjuice 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 Helpjuice. 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	Helpjuice MCP
Server ID	019d75af-81b2-71f3-84b6-ae903ddf52cd
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/helpjuice.