

MCP SERVER

NO CODE

CLOUD HOSTED

# ReadMe MCP

Find any detail across your entire knowledge base.

ReadMe connects your AI client directly to your developer documentation hub. It turns your agent into an expert technical reader, allowing it to instantly find specific guides, API references, or changelogs without leaving the chat window. Your AI can search every page, track recent updates, and pull complex content from your entire knowledge base.

**A+** Quality Score 98.33/100

api-documentation

developer-portal

technical-writing

knowledge-base

documentation-search

developer-experience



# 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](https://cloud.vinkius.com) — connect your AI agent in under 60 seconds.

# ReadMe MCP

10 tools available

Cloud-hosted on Vinkius

This MCP lets you connect your internal developer documentation directly to your agent. You no longer have to copy-paste snippets or switch between tabs just to find an API endpoint definition. Instead, your AI reads the source of truth—your published guides and technical specs—and answers questions using that data.

It handles everything from complex full-text searches across all categories to fetching the exact content for a specific changelog entry. You can even ask it to list all available documentation sections or analyze how different parts of your project are structured. This capability is critical because it keeps your agent grounded in *your* company's official information, not general web data. By connecting this MCP via Vinkius, you ensure that any compatible AI client—whether running locally or on a cloud platform—has immediate access to the full depth and breadth of your technical knowledge.

---

## Core Capabilities

### 01 — Search all documentation

The agent performs full-text searches across every guide, reference page, and custom document.

### 03 — List documentation structure

The agent maps out your entire knowledge base, showing all major categories and sub-guides available.

### 05 — Analyze project structure

The agent provides details on the overall organization and scope of your entire ReadMe project.

### 02 — Retrieve specific content pages

You can fetch the complete Markdown text for any single documented page or API reference.

### 04 — Track product changes

You can pull recent updates, new features, or deprecation notices directly from published changelogs.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/mcp/readme](https://vinkius.com/mcp/readme) — connect your AI agent in three steps.

- 01 Authorize the ReadMe MCP in your environment and provide your API key.
- 02 Instruct your AI client to perform a task, like summarizing an API endpoint or listing all categories.
- 03 The agent uses the tool to read the documentation and responds with accurate, source-backed information.

The bottom line is that you get reliable, real-time access to your entire technical knowledge base from a single conversation thread.

---

## Built For

Technical Writers and Developer Advocates need this. They constantly battle context switching—jumping between their IDE, the docs site, and the chat window—to ensure accuracy. Engineers also benefit when they can instantly pull official API specs into their workflow without leaving their coding environment.

### Technical Writer

Auditing existing documentation for outdated references or finding precise content to draft new guides.

### Developer Advocate

Creating accurate tutorials and quick-start guides by referencing the most up-to-date API definitions and best practices.

### Software Engineer

Quickly pulling official company API documentation into a chat or IDE to validate code against published standards.

---

## What Changes When You Connect

- 01 Instantly check product status: Instead of hunting through multiple sources, use the `list_changeLogs` tool to see a list of all recent updates and announcements immediately.

- 
- 02** Audit content accuracy: If you need to verify an old reference, calling `get_doc` lets your agent pull the full text of that specific page for review, without needing manual logins or searches.
- 
- 03** Understand structure quickly: Use `list_categories` to map out your entire documentation hierarchy. Your agent reports back on all major sections available in a single response.
- 
- 04** Write targeted content: If you need details on one topic but the guide is nested deep, use `get_category_docs` to list all relevant sub-pages before drafting.
- 

---

## Real-World Applications

### **The engineer needs API specs for a new feature.**

Instead of searching Google or asking a colleague, the agent uses `search_docs` and retrieves the exact current endpoint requirements. This lets the software engineer validate their code against official standards in minutes.

### **The technical writer needs to summarize the latest v2 API changes.**

They run `list_changelogs` first, identifying 'v2-api-release'. Then they use `get_changelog` to pull the full content, allowing them to generate a summary for internal stakeholders immediately.

### **The developer advocate needs to create an 'Authentication' guide.**

They ask the agent to list all related documentation (`get_category_docs` for Authentication). The AI provides a comprehensive map, ensuring no necessary sub-guide is missed before writing begins.

### **The team needs to know what's in the project.**

They use `list_custom_pages` and `get_project` together. The agent confirms that key standalone guides exist, providing an overview of content outside the main structured categories.

---

---

# Patterns to Avoid

---

## Asking general questions about APIs

### ✗ AVOID

Asking your AI client, 'How does webhook authentication usually work?' The agent provides generic internet answers that might not match your company's specific implementation.

### ✓ INSTEAD

You must use the `search\_docs` tool. By running a search query through this MCP, you force the agent to ground its answer using only content from your private documentation.

---

## Copying text into the chat manually

### ✗ AVOID

Finding an API reference on one tab and pasting it into the chat window for summarization. This is tedious and fails if the original page changes.

### ✓ INSTEAD

Let the agent use `get\_doc` or `search\_docs`. You simply tell your AI client, 'Summarize the endpoint documented at X,' and it pulls the source content directly.

---

## Assuming documentation is up-to-date

### ✗ AVOID

A colleague points to a guide you found last month. You assume it's current, but it contains deprecated information.

### ✓ INSTEAD

Always check `list\_changelogs` first. This ensures your agent has access to the most recent product updates before trusting any guides.

---

## The Right Fit

Use this MCP if your primary need is for your AI client to answer questions, write drafts, or summarize information based *only* on internal, structured developer documentation (e.g., API specs, user guides, changelogs). If you are building a Retrieval-Augmented Generation (RAG) system that needs deep context from proprietary knowledge bases, this is essential. Don't use it if your goal is to perform general web research or monitor real-time external events; for those, you need generic search tools. Also, don't use this just because the answer sounds nice—always verify that the agent explicitly used a ReadMe tool call (like `get_doc` or `search_docs`) in its response.

---

## Keeping up with documentation is a massive time drain.

Today, getting reliable technical answers requires jumping through hoops. You might search the main site, copy a snippet from an API reference guide, then open the changelog to see if that feature was updated. Then you have to manually paste all those pieces into your chat or document and ask for a summary. It's messy, slow, and highly prone to missing outdated details.

With this MCP connected through Vinkius, the process changes completely. You just talk to your agent. You ask one question—like 'What were the rate limit changes in v2?'—and the agent handles the entire workflow behind the scenes, gathering data from multiple sources and giving you a single, synthesized answer.

---

## ReadMe MCP gives you instant access to your full documentation set.

The manual steps that vanish are the repetitive searches across different tabs. You don't need to manually pull a list of categories, then individually visit each guide to check for related information. The agent organizes it all for you.

Now, your AI client acts as an internal expert who knows every page and every update. It doesn't just point to the documentation; it reads it and answers the question.

---

# ReadMe with 10 Tools

These tools allow your agent to perform every action necessary: searching documents, listing categories, retrieving content, and tracking project updates.

#	TOOL	DESCRIPTION
01	<code>list_categories</code>	Lists all major organizational groups within your entire set of documentation.
02	<code>get_category</code>	Fetches detailed information about one specific section or category of guides.
03	<code>get_category_docs</code>	Lists every single individual documentation page that belongs to a selected guide group.
04	<code>get_doc</code>	Retrieves the full, raw content of any specific documented article or page.
05	<code>search_docs</code>	Runs a comprehensive search across all articles and guides to find relevant information by keyword.
06	<code>list_changelogs</code>	Pulls a list of the titles and dates for all published product update announcements.
07	<code>get_changelog</code>	Retrieves the full text content for one specific, named changelog post.
08	<code>list_custom_pages</code>	Shows a list of standalone pages that exist outside the main category structure.
09	<code>get_custom_page</code>	Retrieves the full content for any specific custom, non-categorized page.
10	<code>get_project</code>	Gathers metadata about the overall scope and configuration of your entire documentation project.

---

## See It in Action

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

### **U** Search the documentation for instructions on configuring webhooks.



Using the `search_docs` tool, I found several results. The top match is 'Webhook Installation Guide' (Slug: `webhook-setup`). It provides step-by-step instructions on setting endpoint URLs.

### **U** Get the contents of the changelog titled 'v2-api-release'.



I requested `get_changeLog` for 'v2-api-release'. The post details new rate limits, updated endpoints for user management, and deprecation notices for the v1 metrics paths.

### **U** List all main documentation categories.



I fetched the configuration using `list_categories`. There are 4 global categories: 'Getting Started', 'Authentication', 'REST API Reference', and 'Webhooks Guide'.

---

## Frequently Asked Questions

### **01** How does ReadMe MCP help me find a specific API endpoint?

You use the `search_docs` tool. Instead of guessing keywords, your agent performs a full-text search across all guides and returns direct links and content snippets matching the endpoint you need.

### **02** Can I list everything in my documentation?

Yes, start by using `list_categories` to see the main sections. Then, use `get_category_docs` on any category name to get a full manifest of all related pages.

---

**03 How do I know what was changed recently?**

Run the `list\_changelogs` tool first. This gives you a list of titles and dates for all updates, helping you pinpoint exactly which version you need to read about next with `get\_changelog`.

---

**04 Does ReadMe MCP work only on guides?**

No, it's more comprehensive. Besides structured guides, you can also check standalone content by using `list\_custom\_pages` and fetching the details with `get\_custom\_page`.

---

**05 What if I need to summarize a whole guide?**

You should use `get\_doc` on the specific page you want summarized. While the tool fetches the content, your agent can then process that full text into the summary you request.







---

# Go Live in 60 Seconds

Get your connection token from [cloud.vinkius.com](https://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
 <b>Claude AI</b>	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 <b>Cursor</b>	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 <b>VS Code</b>	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"readme": { "url": "..." }</code>
 <b>Windsurf</b>	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 <b>ChatGPT</b>	Settings → Tools & plugins → Add MCP server → Paste endpoint
 <b>Gemini</b>	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

# ReadMe 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](https://vinkius.com) · [support@vinkius.com](mailto: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 ReadMe. 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	ReadMe MCP
Server ID	019d75fc-ed18-7242-997a-4e24253948d1
Platform	Vinkius Cloud for AI Agents
Endpoint	<a href="https://edge.vinkius.com/{token}/mcp">https://edge.vinkius.com/{token}/mcp</a>

### 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/readme](https://vinkius.com/mcp/readme).