MCP SERVER | NO CODE | CLOUD HOSTED

Tettra MCP

Access your entire company knowledge base from anywhere.

Tettra MCP connects your internal wiki documentation directly to your AI agent. Need an API spec or a team policy? Your agent searches, reads, and even drafts new content from across your company's knowledge base, all without you leaving your IDE. It handles full-text searching, page retrieval, and publishing updates instantly.

A+ Quality Score 100/100

wiki | internal-documentation | search | knowledge-base | team-collaboration | onboarding

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

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.

Tettra MCP

12 tools available
Cloud-hosted on Vinkius

Stop switching tabs just to look up internal details. This MCP connects Tettra, your company's central documentation hub, right into your developer workflow. Instead of clicking through multiple categories or running separate searches, you talk to your agent and it handles the whole process. You can ask complex questions and get instant answers by searching across all existing pages. Need to contribute? Your agent drafts new wiki content or updates old guides on demand. This means keeping your team's knowledge current is part of your daily coding routine.

When you connect this MCP through Vinkius, you give your AI client a direct line into your organizational memory. It's more than just searching; it lets your agent publish new documentation and even manage the Q&A board, making sure your team always has accurate information at their fingertips.

Core Capabilities

01 — Search across all wiki pages

Find instant answers by performing a full-text search across every page in your Tettra knowledge base.

02 — Retrieve full documentation content

Read the complete text, including markdown and HTML, for any specific technical guide or policy page inside your chat window.

03 — Create new wiki pages

Draft and publish brand-new internal articles in a specified category with simple prompts.

04 — Update existing documentation

Modify the title or content of an old page using its ID, keeping your guides current.

05 — Manage Q&A records

List existing questions or post new ones directly to the team's internal question board.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and enter your Tettra Team ID and API Key (requires a Professional or Scaling plan).

02 Instruct your AI client to perform an action, like listing categories or searching for a term.

03 The agent executes the request against Tettra, returning formatted data—whether it's page content, a list of questions, or confirmation that a new article was published.

The bottom line is you use your natural language prompts to get your AI agent to manage and query all aspects of your company knowledge base.

Built For

Engineers, Ops Leads, and Product Managers who spend too much time context-switching between code editors, documentation sites, and Slack threads. You need instant access to the 'how' and 'why' behind every piece of software.

Software Engineer

Uses this MCP to look up architecture decisions or environment setup steps without leaving their IDE.

Operations Lead

Rapidly generates and publishes standard operating procedures (SOPs) using AI prompts, keeping the team manual current.

Technical Writer/Product Manager

Verifies outdated documentation pages or suggests new articles via chat to maintain a robust knowledge base.

What Changes When You Connect

01 Instant Answers: Use the search_pages tool to find specific answers across all documentation without running manual keyword searches or browsing folder trees.

02 Zero Context Switching: Your agent reads full page content via get_page_content, providing technical guides and policies right inside your chat window while you code.

03 Documentation Maintenance: Easily publish new knowledge using create_wiki_page. This keeps internal processes (like SOPs) immediately available to every team member.

04 Knowledge Governance: Use verify_wiki_page to mark critical documents as up-to-date, building trust and reliability in your company's source of truth.

05 Systematic Q&A Management: Instead of checking a separate board, list_qa_questions lets you see unanswered questions or post new ones right from your workflow.

Real-World Applications

Debugging an unknown API requirement

An engineer needs to know the exact format for a specific payload. They ask their agent, and it uses search_pages and get_page_content to pull up the relevant section from the 'API Specs' wiki page, providing the necessary data instantly.

Creating an SOP after a process change

An operations lead finishes updating a workflow. They prompt their agent to create_wiki_page in the 'Support' category and draft the new Standard Operating Procedure, which is then immediately visible on Tettra.

Identifying necessary policy updates

A product manager notices inconsistencies across documentation. They ask their agent to list_pages_in_category for 'Legal' and identify pages that haven't been updated in six months, flagging them for review.

Addressing unanswered team questions

A developer runs into a recurring roadblock. They use the agent to list_qa_questions, see that the question is still open, and then use create_qa_question to formally post it to get an official answer.

Patterns to Avoid

Assuming general search works

✗ AVOID

Trying to find a specific API version detail by just searching Google for 'Tettra API v2'. This returns external, unverified results that might be wrong.

✓ INSTEAD

Use the Tettra MCP's built-in tools. Tell your agent to search_pages within the knowledge base or get_page_content of a specific guide ID.

Manually copying content

✗ AVOID

A user reads an old policy page, copies the text into Slack, and pastes it for discussion. This loses formatting and makes version control impossible.

✓ INSTEAD

Use get_page_content to retrieve the full markdown/HTML body directly inside your chat client. The agent keeps the original context and structure.

Forgetting documentation updates

✗ AVOID

A team member updates a process but forgets to mark it as current, leading others to use outdated steps.

✓ INSTEAD

After publishing new content or updating an old one with update_wiki_page, immediately run verify_wiki_page on that page ID. This signals the system and the team that it's accurate.

The Right Fit

Use this MCP if your core bottleneck is finding, creating, or maintaining structured institutional knowledge stored in a wiki format (Tettra). You need an agent that acts as a 'knowledge librarian,' retrieving specific articles and helping you publish them. Don't use it if you are trying to connect Tettra to live operational systems (like Jira tickets or GitHub repos) — for those, look at integration tools designed specifically for workflow automation. If your goal is just general internet search or unstructured data analysis, a standard web retrieval tool will suffice. But when the knowledge *is* the system, this MCP is essential.

The problem with internal documentation today isn't bad writing; it’s friction.

Right now, finding a basic procedure requires navigating three different tabs: first to find the correct wiki category, then opening the page, and finally copying the relevant snippet into Slack so someone else can see it. This constant copy-pasting across platforms slows down every decision.

With this MCP, you simply ask your agent for the required information—for example, 'What is the payout calculation?'—and it retrieves the full, properly formatted content from Tettra directly in your chat. You get instant context and zero tab switching.

Tettra MCP: Publish and update knowledge without leaving your agent.

Before this tool, drafting a new Standard Operating Procedure meant multiple steps: writing the content in a separate editor, submitting it for review, waiting for approval, and then manually publishing it to the correct wiki folder. It was slow and required coordination.

Now, your agent handles the heavy lifting. You use create_wiki_page to draft and publish brand-new articles instantly, or you run update_wiki_page when a policy changes, keeping your documentation accurate the second the decision is made.

Tettra: 12 Tools for Documentation Ops

These tools let your agent search, read, write, and manage every aspect of your company's Tettra wiki and Q&A system.

#	TOOL	DESCRIPTION
01	`create_wiki_page`	Creates a brand-new wiki page within a specific category using a title, content, and category ID.
02	`create_qa_question`	Posts a new question directly into the Tettra Q&A system for team members to answer.
03	`get_category_details`	Retrieves specific metadata and details about a designated Tettra category.
04	`get_page_content`	Fetches the full title, markdown, and HTML body content of an existing Tettra page.
05	`list_categories`	Retrieves a list of all top-level categories available across the entire Tettra wiki.
06	`list_pages_in_category`	Returns a list of all specific wiki pages that belong under one selected category.
07	`list_qa_questions`	Retrieves a comprehensive list of all questions currently posted in the Tettra Q&A system.
08	`list_subcategories`	Lists all subfolders or subsections nested under a specific parent category.
09	`search_pages`	Performs a full-text search across the entire Tettra wiki and returns up to five matching pages.
10	`suggest_new_page`	Automatically suggests that a new wiki page needs to be created for the team's review.
11	`update_wiki_page`	Modifies the title or content of an existing Tettra page using its unique ID and new information.
12	`verify_wiki_page`	Marks a specific wiki page as verified, confirming it is up-to-date and accurate.

See It in Action

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

U Search the wiki for 'Database Migration Checklist'.

I found 1 relevant page in the wiki:
- **Database Migration Checklist** (ID: 1042) in the 'Engineering' category.
Would you like me to fetch the full markdown content of this guide?

U Create a new wiki page in the 'Support' category explaining how to handle refund requests.

Successfully published page **Refund Handling Procedure** (ID: 4055) to the Support category (ID: 12)!
The content has been formatted with step-by-step markdown. Your team can now search for it natively on Tettra.

U Mark page ID 883 as verified and up to date.

The page 'Deployment Keys Matrix' (ID 883) has been officially verified! ✅
The system will log that it was verified today, maintaining trust in your knowledge base.

Frequently Asked Questions

01 **How do I search for specific guides using Tettra MCP?**

You use the agent to run search_pages. You just give it a topic—like 'Database Migration Checklist'—and the agent returns links and titles from your knowledge base.

02 **Can I publish new wiki pages using Tettra MCP?**

Yes, you can use create_wiki_page. You provide a title, content, and category ID, and the agent publishes the article to Tettra for your team.

03 **What if I update an old page? Do I need multiple steps?**

No. If you know the page ID, you use update_wiki_page with the new content. To make sure everyone trusts it afterward, run verify_wiki_page.

04 **How does Tettra MCP help with Q&A management?**

You can list_qa_questions to see unanswered topics or create_qa_question to formally submit a new question for the team to address.

05 **Does this connector work in VS Code? (Tettra MCP)**

Yes, it works across any MCP-compatible client. You connect your preferred AI agent inside Cursor or VS Code and use Tettra's tools just as normal.

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 `"tettra": { "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

Tettra 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 Tettra. 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	Tettra MCP
Server ID	`019d7612-5594-72a6-ba57-479c60cd0835`
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/tettra.