MCP SERVER | NO CODE | CLOUD HOSTED

Cody AI MCP for AI Agents

Manage proprietary knowledge base documents and support bots

Cody AI connects your agent to a knowledge base trained on proprietary documents and web content. It lets your AI client manage specialized bots, import new company guidelines from URLs or files, and start conversations using only your private information.

A+ Quality Score 100/100

rag | ai-assistant | knowledge-base | chatbot-training | document-query | automated-support

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/cody-ai

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.

Cody AI MCP

10 tools available
Cloud-hosted on Vinkius

When customer questions are scattered across wikis, PDFs, and old support tickets, keeping an AI agent accurate is a nightmare. This MCP solves that by connecting to Cody AI, turning your internal documentation into actionable knowledge for your agents. You feed the system your corporate manuals, web pages, and policy documents, and it creates specialized bots trained only on that data. Your agent doesn't guess; it pulls precise answers directly from your source material.

Through Vinkius, you connect this capability to any compatible AI client—whether it's Claude or Cursor. You can manage the entire lifecycle of your knowledge base right through natural conversation. Need to add a new HR policy? Simply import that web page and let the system handle the rest. The bottom line is, your team gets instant, accurate answers from company data without ever needing manual API calls.

Core Capabilities

01 — Train and List Knowledge Bots

Retrieve a list of all configured bots and fetch detailed information about any specific bot to check its current training status.

02 — Manage Knowledge Sources

Import content from external URLs or retrieve lists of existing folders and documents within your knowledge base structure.

03 — Monitor Document Syncing

Check the syncing status of an imported document to confirm that the AI has finished processing it and is ready to answer questions based on its contents.

One Click on Vinkius — From Prompt to Execution

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

01 Add the Cody AI integration to your preferred AI client's toolset.

02 Provide your API Key from the Cody AI dashboard. Your agent now recognizes the knowledge base connection.

03 Use natural language prompts (e.g., 'Show me all bots') to manage bot configurations, import content, and start conversations.

The bottom line is that you treat complex data management—like importing a web page or checking document status—just like chatting with your agent.

Built For

This MCP is for operations teams, support managers, and knowledge architects. If your job involves turning scattered documents into reliable answers, this is for you. It helps move beyond basic chatbots toward true corporate intelligence.

Support Team Lead

Manages instant access to customer questions by querying trained bots using the latest product manuals or service agreements.

Knowledge Manager

Updates the central knowledge base by importing new web pages and monitoring document syncing status from chat, eliminating manual content ingestion processes.

What Changes When You Connect

01 Stop copy-pasting answers. Use the `send_message` tool to query a bot directly against your company's entire document library, giving instant, source-cited answers.

02 Never start from scratch. Run `list_bots` to see every available specialized agent—from HR policy bots to product FAQ bots—in one view.

03 Keep documentation current. Use `import_webpage` to feed new guidelines or policies directly into the knowledge base via a simple URL, keeping your agents up-to-date automatically.

04 Stay organized with `list_folders`. You can see exactly where every document lives and group related content before training a bot on it.

05 Know when data is ready. Use `get_document_status` to confirm that newly uploaded files have finished syncing, so your agent doesn't answer questions based on incomplete knowledge.

Real-World Applications

Onboarding a new product line.

A Product Manager needs the AI to answer niche questions about a new feature. Instead of manually uploading dozens of PDFs, they use `import_webpage` to feed all the launch site documentation into a bot's knowledge base, making it instantly available for agent querying.

Handling complex compliance inquiries.

A support agent gets asked about conflicting policies. They ask their AI client to check multiple bots and use `list_messages` to review the full chat history, ensuring they cite the correct policy document from the knowledge base.

Creating a dedicated department bot.

A Knowledge Manager wants an HR bot that only talks about employee guidelines. They first use `list_bots` to confirm their current setup, then create and configure a new bot using `get_bot_details`, restricting its scope to the relevant folders.

Troubleshooting outdated guides.

A team notices an old support bot is referencing deprecated procedures. They use `list_documents` to find the original source file and then run `get_document_status` on the replacement guide, ensuring the new data has successfully replaced the old knowledge.

Patterns to Avoid

Building a bot with random files.

✗ AVOID

Just uploading 100 PDFs into one place and letting the agent talk to them. This leads to vague, contradictory answers because there's no structure or source control.

✓ INSTEAD

First, use `list_folders` to organize your content by department (e.g., HR, Billing). Then, train separate bots on those specific folders. If you need a new policy, import the single web page using `import_webpage` before linking it to any bot.

Assuming knowledge is instant.

✗ AVOID

Importing 50 large documents and immediately asking the AI questions, only for the answers to be vague or wrong because the system hasn't finished indexing everything.

✓ INSTEAD

Always check the `get_document_status` tool after a batch upload. Wait until the status confirms 'complete' before routing any critical queries through that bot.

Starting over every time.

✗ AVOID

Every time you need to talk to a specific bot, having to re-enter all context and ask basic setup questions.

✓ INSTEAD

Use `list_bots` to identify the correct agent name. Then, use natural language commands to create a new session with the targeted bot using `create_conversation`, saving you time.

The Right Fit

Use this MCP if your biggest headache is turning siloed company documents—whether in PDFs, wikis, or URLs—into reliable answers for support staff. You need an agent that *knows* your internal policies and can cite where it found the information.

Don't use it if you just need a generic chatbot that pulls from general web knowledge (like Wikipedia). For those needs, other specialized chatbots are better suited. Also, don't try to manage complex user permissions; this MCP focuses purely on document ingestion, bot configuration (`get_bot_details`), and conversation flow. If your primary job is managing user accounts or access rights, you need an identity management tool instead.

Cody AI for Knowledge Base Management: Centralizing Support Docs

Right now, answering a customer question means checking five different places: the Jira ticket, the internal SharePoint site, the PDF manual, and the latest web page. Your support team spends half their day copy-pasting snippets and cross-referencing which document is actually current.

With this MCP, your agent does it for you. You use `import_webpage` to pull new guidelines from a single URL, and the bot incorporates that data immediately. The result? Your agent gives one authoritative answer, citing exactly where in your knowledge base it found the policy.

Cody AI for Document Workflow: Structuring Bot Training Data

Manual bot training involves tedious steps: exporting data into chunks, cleaning up redundancy, and manually feeding those chunks to a separate model. This process is slow and requires dedicated technical staff just for maintenance.

This MCP simplifies the whole workflow. You use `list_folders` to structure your source materials first, then you can train bots on specific groups of documents. It moves documentation management from a manual engineering task to a simple data flow operation.

Cody AI MCP: 10 Tools for Knowledge Base Bot Operations

These tools let you manage, update, and interact with all the specialized bots that use your internal documents as their knowledge source.

#	TOOL	DESCRIPTION
01	`create_conversation`	Starts a fresh conversation session with a specific bot for immediate Q&A.
02	`get_bot_details`	Retrieves comprehensive details about an existing knowledge bot, including its training scope and configuration.
03	`get_document_status`	Checks if the AI has finished processing a document so you know exactly when it's ready for questions.
04	`import_webpage`	Takes content from a specific URL and adds it to your knowledge base, expanding a bot's data set.
05	`list_bots`	Shows every single specialized bot currently active in the Cody AI account.
06	`list_conversations`	Gives you a list of your most recent chat sessions for quick review.
07	`list_documents`	Retrieves an organized list of all documents currently stored in the knowledge base.
08	`list_folders`	Shows the folder structure of your entire knowledge base, helping you organize content logically.
09	`list_messages`	Retrieves the full message history for a specific conversation thread.
10	`send_message`	Sends a direct prompt or question to the AI within an active conversation.

See It in Action

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

U List all the specialized bots we have right now.

Knowledge Base Bots Overview

You currently run 4 active knowledge bots:

- 🧑‍🤝‍🧑 **HR Policy Bot:** Trained on Employee Handbook (2023) and Benefits Guide.
- ⚙️ **Product FAQ Bot:** Focuses on setup guides and troubleshooting.
- 💰 **Billing Bot:** Handles invoice questions and payment plans.
- 📜 **Legal Compliance Bot:** Uses the latest regulatory documents.

Which bot should we start a conversation with?

U I found this new article about our updated privacy policy, can you add it to the Legal Bot?

✅ Upload initiated for 'new_privacy_policy.html'.

The system is now processing the content and indexing new vectors. The bot will be fully ready to answer questions regarding data retention in approximately 3 minutes.

Source files: `legal_dept_folder/`
Current status: Indexing (75% complete)

U Show me the full chat history for the Product FAQ Bot regarding Widget X.

Conversation History: Product FAQ Bot

DATE	USER QUERY	BOT RESPONSE SUMMARY
2024-07-15	How do I reset my widget?	Go to Settings > Profile. (Source: *Setup Guide v2*)
2024-07-16	What's the max throughput?	The limit is 50 units/hour. (Source: *Technical Spec Sheet*)

Need to send a new message or start a new topic?

Frequently Asked Questions

01 **How does Cody AI MCP help me update my company policies?**

You can keep your knowledge base current by using the import features. You simply provide a URL, and the system automatically pulls that web page content into the correct bot's training data so it has the latest information.

02 **Is Cody AI MCP better than just uploading PDFs to my chatbot?**

Yes. This MCP lets you manage the whole process, including knowing when a document is ready (`get_document_status`). You can also organize content into specific bots for different topics, preventing confusion from mixed documents.

03 **Can I use Cody AI MCP to train a bot on my internal wikis?**

Yes. If your wiki pages are available online via a URL, you can feed them directly using the `import_webpage` tool. This is much faster and more reliable than manual data handling.

04 **What if I need an agent to talk about multiple topics?**

You should create separate bots for different topics (like HR, Billing, etc.). The MCP helps you manage these distinct agents using `list_bots` and ensures each one only uses the data relevant to its purpose.

05 **Does Cody AI MCP handle document organization?**

It does. You can use tools like `list_folders` to see your knowledge base structure, ensuring that when you add new content, it lands in the right place before training a bot on it.

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 `"cody-ai": { "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

Cody AI 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 Cody AI. 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	Cody AI MCP
Server ID	`019d7576-b8d1-7185-b492-786aa5545bc6`
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/cody-ai.