MCP SERVER | NO CODE | CLOUD HOSTED

Nuclino MCP

Manage all knowledge and project documentation.

Nuclino MCP gives your AI agent full read/write access to corporate knowledge bases, team structures, and project documentation. It treats your entire wiki—including nested workspaces, collections, and files—as one searchable source. You can query company history, draft new specs, or track user roles directly inside your preferred client.

A+ Quality Score 100/100

team-wiki | document-collaboration | workspace-management | knowledge-base | project-tracking

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

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.

Nuclino MCP

12 tools available
Cloud-hosted on Vinkius

Connect your AI agent to your internal knowledge base for reading, writing, and organizing corporate data. This MCP lets you interact with your company's entire workspace structure—everything from team listings down through individual documents. You can ask the agent to find a specific policy across multiple departments or draft an update directly into a living project page. It tracks documentation evolution over time, making sure any new guideline or meeting note gets recorded instantly. Through Vinkius, you get one connection point that lets your AI client access this deep organizational data without needing dozens of separate API keys.

Core Capabilities

01 — Map the entire team structure

Discover every root team and nested workspace ID within your organization's hierarchy.

02 — Find specific documents globally

Run a fuzzy search across all knowledge items to pinpoint exact policies, notes, or project pages no matter where they live.

03 — Read and analyze content

Retrieve the full Markdown payload of any item, along with listing all associated physical files and attachments.

04 — Draft new documentation

Instantly create a brand new knowledge page or document within an existing workspace using natural language prompts.

05 — Manage structure and users

View who is attached to the team, map customizable fields, or list all collections that group related documents.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and enter your Nuclino Personal API Key.

02 Select a task in your AI client (Claude, Cursor, etc.), telling the agent what knowledge you need or want to create.

03 The agent uses its tools—like `search_items` or `create_item`—to interact with your wiki and deliver the result.

The bottom line is that your AI agent treats your corporate wiki like a database, giving you structure and content access all in one go.

Built For

This MCP is for anyone whose job depends on institutional memory: the tech writer who spends hours cross-referencing old guidelines, the project manager struggling to keep track of scattered deliverable specs, or the engineer needing immediate access to outdated API documentation.

Technical Writer

Drafting new sections requires finding existing policy documents using `search_items`, reading their contents via `get_item` for accuracy, and then appending the new text.

Project Manager

Need to kick off a new initiative? Use the agent to create the main planning document with `create_item`, identify necessary stakeholders using `list_users`, and map out the scope by listing collections.

Software Engineer

When tackling a bug, you use the agent to search technical specs across departments via `search_items` or list files (`list_files`) attached to an existing feature item for context.

What Changes When You Connect

01 Stop digging through nested folders. Use `search_items` to run a global query across your entire team's wiki, finding policies or specs regardless of which workspace they live in.

02 Drafting updates used to mean copying old text into a new doc. Now, you can use the agent to read existing content via `get_item`, draft the changes, and then update it using `update_item`—all in one go.

03 You don't have to manually track who owns what. The MCP lets you list all team members (`list_users`) and map out ownership fields (`list_fields`), keeping your documentation accurate automatically.

04 Need a new planning document? Instead of creating it manually, the agent can use `create_item` instantly, generating a live, structured page that awaits content input.

05 The MCP exposes structural details like collections and workspaces. You get visibility into how your knowledge base is organized via tools like `list_collections`, which saves hours of mapping time.

Real-World Applications

Onboarding a new team member

A new PM needs to understand the department's process flow. Instead of asking three different people for links, they ask their agent to search all documentation mentioning 'onboarding workflow'. The agent runs `search_items` and compiles the top five relevant pages.

Updating a critical policy

The compliance team needs to revise an old security guideline. They tell their agent, which uses `get_item` to pull the existing Markdown content. The agent drafts the revision and then executes `update_item`, ensuring the change is live immediately.

Finding a specific attachment

The engineer needs the final PDF spec sheet for Project Chimera, but they only know it was attached to 'Project Alpha'. They use the agent to find the item and list all files using `list_files` to get the direct download link.

Structuring a new project

A department is launching a massive initiative. The PM asks their agent to create the main planning hub, which uses `create_item`. They then ask the agent to list all available team members using `list_users` to set up initial roles.

Patterns to Avoid

Searching by vague keywords

✗ AVOID

Just typing 'security' into a general search box and getting hundreds of irrelevant results, forcing you to manually click through dozens of pages.

✓ INSTEAD

Use the agent with `search_items` and specify the context: 'Search for security policies related to API access in the Engineering team.' This focuses the scope dramatically.

Assuming content is fully updated

✗ AVOID

Relying on an old wiki page that hasn't been manually edited since last quarter, meaning key process steps are missing or incorrect.

✓ INSTEAD

Always verify core documents. Use `get_item` to pull the raw Markdown payload and review it directly in your agent client before trusting its contents.

Creating content without structure

✗ AVOID

Pasting a block of notes into a random spot, making it hard to find or relate to other documentation.

✓ INSTEAD

When drafting, use `create_item` and immediately ask the agent to list available fields (`list_fields`). This forces you to categorize the data correctly from the start.

The Right Fit

Use this MCP if your primary pain point is managing institutional memory within a structured wiki environment. If your team's knowledge lives in Nuclino, and you need an AI agent to read, write, or organize that content—this is it. You get deep document access via `get_item` and structural visibility using `list_collections`. Don't use this if you are managing unstructured data outside of the wiki (like emails, Slack threads, or external CRM records). For those

needs, look for a messaging MCP or a dedicated CRM connector instead.

Finding institutional knowledge feels like being an archaeologist.

Today, finding the single source of truth is painful. You don't know if the latest compliance guidelines are buried in the '2023 Marketing Docs' workspace or if they were uploaded as a file attachment to the old 'Project Phoenix' item. You end up clicking through three different department sites, copying fragments into a meeting summary, and hoping you didn't miss anything critical.

With this MCP, your agent performs global searches using `search_items`. You simply ask it: 'What is the current policy for international data transfer?' It pulls from every relevant corner of your wiki instantly. It doesn't just find a link; it gives you the context.

The Nuclino MCP empowers you to write, not just read.

Before this, drafting a new project plan meant manually creating the item, ensuring the right people were assigned, and then hoping someone else would update it with the final specs. It was a multi-step handoff process full of lag time.

Now, you tell your agent to create the main document using `create_item` in the correct workspace. Then, you can ask it to gather all relevant user details via `list_users`, ensuring every necessary owner is assigned from the jump. The entire documentation lifecycle moves at the speed of conversation.

Nuclino MCP: 12 Tools for Knowledge Management

These tools give your agent full control over listing, reading, creating, updating, and searching every structural element inside a Nuclino workspace.

#	TOOL	DESCRIPTION
01	`list_teams`	Lists all organizational teams the authenticated user belongs to, helping locate specific departments within your company structure.
02	`list_workspaces`	Returns a list of isolated workspaces mapped within a specified team, allowing you to scope subsequent queries accurately.
03	`list_items`	Lists all standard knowledge items (pages) in a workspace, giving you an overview of the available documentation titles.
04	`get_item`	Retrieves the precise Markdown payload and full configuration details for a specific document item.
05	`create_item`	Writes a brand new knowledge page or wiki item into a specified workspace, making it immediately visible to your team.
06	`update_item`	Appends new edits or overwrites partial content inside an existing documented item.
07	`delete_item`	Permanently removes a structural knowledge item from the wiki after confirming with you.
08	`search_items`	Executes a global semantic search across an entire team's content to find unknown documents or policies by keyword.
09	`list_collections`	Lists the document collections that segment and organize a target workspace, helping you understand the document relationship graph paths.
10	`list_fields`	Maps customizable structured property fields available globally across an entire team's knowledge base.
11	`list_users`	Lists all human identities attached to a specific team, giving you a roster of active members and owners.
12	`list_files`	Retrieves the URLs for physical attachments bolted onto an item, allowing you to access source files like PDFs or images.

See It in Action

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

U Search Nuclino for any documentation mentioning 'SSO Security Policies'.

I've successfully performed a fuzzy text search and located the page 'SSO Security Implementation Guidelines'. I can extract its contents or share the browser link if needed.

U Create an item titled 'Project X Architecture Brief' in the Engineering workspace.

I retrieved the Engineering workspace ID and created the new item 'Project X Architecture Brief'. The document is live and awaiting further edits.

U List all teams connected to this authentication token.

You have access to 3 teams on this environment: Design Ops (`id-1`), Software Engineering (`id-2`), and Executive Planning (`id-3`).

Frequently Asked Questions

01 **How does Nuclino MCP search for documents?**

The agent performs indexed semantic searches across an entire team using `search_items`. This means it doesn't just match keywords; it understands the context of your query to find related content.

02 **Can I use Nuclino MCP to update a document?**

Yes, you can. The agent uses `update_item` to append new text or overwrite partial states in an existing wiki page, keeping your documentation current without manual copy-pasting.

03 Does Nuclino MCP only list items by title?

No. Beyond listing titles with `list_items`, you can read the exact content and configuration of any item using `get_item` to see the full Markdown payload, including hidden metadata.

04 What if I need to find a file attached to an item?

You use the `list_files` tool. It exposes pure URL bindings for physical attachments on any specific document item, letting you access source files like PDFs or images.

05 How do I start using Nuclino MCP in my AI client?

You subscribe to this MCP on Vinkius and provide your Nuclino Personal API Key. After that, you just tell your agent what documentation task needs completing.

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 `"nuclino": { "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

Nuclino 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 Nuclino. 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	Nuclino MCP
Server ID	`019d75e0-0a09-7232-ae3b-31cbe629698b`
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/nuclino.