MCP SERVER | NO CODE | CLOUD HOSTED

Nyckel ML MCP

Classify Data and Find Patterns with AI Chat

Nyckel ML connects your AI agent to advanced machine learning tools for automated data classification and semantic search. You can test custom models, classify text or images instantly, and find similar samples using natural language—all without writing a single line of integration code. It lets you manage the entire lifecycle of your ML assets right from your chat client.

A+ Quality Score 100/100

machine-learning | classification | semantic-search | automated-labeling | predictive-modeling | data-tagging

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

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.

Nyckel ML MCP

10 tools available
Cloud-hosted on Vinkius

This MCP gives your AI agent access to professional machine learning capabilities. Instead of building complex APIs or running batch jobs, you simply ask your agent to classify data or search a gallery. Need to know if an uploaded image is a product or just clutter? You prompt the system, and it runs the appropriate ML function, giving you instant predictions along with confidence scores. It's useful for everything from content moderation to e-commerce research.

If your team needs to build custom data workflows, this connection makes it possible. Your agent can list existing functions or look at training samples to check accuracy before making a prediction. When you connect this MCP via Vinkius, you get access to all these features through one conversational point. You're doing deep ML work, but the interaction feels like just asking a smart teammate for an opinion.

Core Capabilities

01 — Classify Content

Send text or image URLs and receive instant predictions and confidence scores from your pre-trained machine learning functions.

02 — Perform Semantic Search

Query existing search galleries to find samples that are conceptually similar, even if they don't contain the same keywords.

03 — Manage ML Functions

List and retrieve detailed metadata for all machine learning functions available in your Nyckel account.

04 — Curate Training Data

Upload new training samples, assign labels, or delete entire ML functions to refine model performance.

05 — Check Account Status

Retrieve profile and workspace metadata for the authenticated Nyckel account.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and enter your unique Nyckel Client ID and Secret credentials.

02 Your AI client connects, establishing a secure link that grants access to all ML tools.

03 You issue a natural language command through your agent—for example, 'Classify this text' or 'Find samples similar to this image.'—and the MCP executes the required function.

The bottom line is you talk to your AI agent like normal, and it handles all the complex data processing in the background.

Built For

This connector serves Data Scientists who need to validate model outputs before deployment. It's perfect for Content Moderators drowning in user-generated content and AI Developers rapidly prototyping ML ideas without custom coding.

Content Moderator

Automates the classification of massive volumes of user-submitted text or images, ensuring consistent labeling across all incoming data.

Data Scientist

Monitors training samples and prediction accuracy in real time. They can use tools like list_ml_samples to review data quality or annotate_ml_sample to manually correct labels.

AI Developer

Tests classification models or search galleries instantly by invoking_ml_function, skipping the step of writing boilerplate API integration code.

What Changes When You Connect

01 You get immediate, actionable predictions. Instead of waiting for a batch job or writing custom code to hit an endpoint, you simply ask your agent to run the classification via invoke_ml_function.

02 Your search is smarter. Using semantic_search means you don't have to guess keywords; your agent finds samples that are conceptually related to what you provide.

03 Data governance becomes easy. You can list all available labels using list_ml_labels, ensuring your classification process sticks to the defined schema every time.

04 You stay in control of your data pipeline. The MCP lets you monitor training progress by listing_ml_samples and manually assigning or updating tags with annotate_ml_sample.

05 Rapid prototyping is possible. AI developers can test multiple ML functions by list_ml_functions without ever leaving their chat environment.

Real-World Applications

Automating Content Screening

A content moderator receives a flood of user messages and needs to classify sentiment and detect prohibited imagery. They simply tell their agent, 'Classify these 50 images using the Sentiment Classifier.' The agent executes invoke_ml_function for each image and returns a summary report with confidence scores.

Debugging Model Performance

A data scientist suspects one of their ML functions is biased. They use list_ml_samples to pull up the raw training data, then manually annotate_ml_sample on 20 records to check if human input aligns with the model's current labels.

Finding Product Inspiration

An e-commerce designer uploads a sketch of a new product and needs to see similar items sold previously. They prompt their agent, 'Find me products like this drawing,' triggering semantic_search against the entire product gallery.

Checking Model Scope

A developer joins a project mid-cycle and doesn't know what ML tools exist. They ask their agent to list all available functions using list_ml_functions, getting an instant overview of the entire system.

Patterns to Avoid

Treating it like a simple database query

✗ AVOID

Trying to find data by exact text match when all you really need is conceptual similarity. You'd ask, 'Show me everything about blue shoes,' but the search results are too narrow.

✓ INSTEAD

Don't use simple keyword queries. Use semantic_search instead. This tool finds samples that *mean* 'blue shoes' even if they are labeled 'indigo footwear'.

Ignoring data governance rules

✗ AVOID

Manually changing a model's definition in the UI without knowing which labels are valid, leading to inconsistent classification results.

✓ INSTEAD

Before modifying anything, always call list_ml_labels. This guarantees you know exactly what categories and tags your ML models expect.

Overwriting training data accidentally

✗ AVOID

Mistakenly labeling a clean, accurate sample as incorrect or deleting an entire function because it wasn't working on the first try.

✓ INSTEAD

Always review list_ml_samples and get_ml_function before making changes. Use annotate_ml_sample only after confirming the source data is correct.

The Right Fit

Use this MCP if your workflow requires making decisions based on machine learning predictions, whether classifying content or searching deep knowledge bases. The power here lies in running complex ML functions—like invoke_ml_function—using simple natural language prompts, keeping you inside your chat agent experience. Don't use it if your goal is merely data storage; for that, connect a dedicated database connector. If all you need to do is manage basic user records or send emails, look for messaging or CRM-type MCPs instead. This tool is specialized for the full ML lifecycle: from creating samples with create_ml_sample, through prediction via invoke_ml_function, and finally monitoring everything using list_ml_functions.

Handling data classification used to be a nightmare of clicks.

Today, if you had to classify thousands of pieces of user-generated content—say, categorizing customer feedback or screening images for policy violations—you'd spend hours in a dedicated portal. You'd have to manually upload the file, select the correct ML model from a dropdown list, input specific parameters (like confidence thresholds), and then hit 'Run,' only to wait for the results to populate a separate dashboard that you had to cross-reference.

With this MCP, those steps disappear. Instead of navigating multiple tabs or dealing with complex forms, you just ask your agent, 'Classify these 50 images using the Product Search function.' The agent handles the entire sequence—calling the right tool, passing the data, and getting the clean results back to you in plain text.

Nyckel ML MCP gives you control over your models.

Before this connector, monitoring a model meant logging into the platform's internal dashboard. You'd have to manually check which samples were used for training and if any labels needed updating by comparing raw data against the current metadata set.

Now, you can ask your agent to list_ml_samples or get_ml_function details directly. It gives you full visibility into your ML assets without ever leaving your chat window.

Nyckel ML: 10 Tools for ML Management

These ten tools allow you to fully manage the lifecycle of your machine learning assets, from creating samples to running live predictions.

#	TOOL	DESCRIPTION
01	`annotate_ml_sample`	Assigns a specific label to an existing data sample in your training set.
02	`create_ml_sample`	Adds a brand new piece of raw data to be used as a training sample for your models.
03	`delete_ml_function`	Permanently removes an existing machine learning function from your account.
04	`get_ml_function`	Retrieves specific configuration and metadata details for a single ML function by its ID.
05	`get_account_info`	Fetches general profile and workspace information about your connected Nyckel account.
06	`invoke_ml_function`	Runs a specific, trained ML function against new data to get an instant classification or prediction score.
07	`list_ml_functions`	Lists all the machine learning functions currently defined within your account.
08	`list_ml_labels`	Retrieves a comprehensive list of every available label and category used by your ML models.
09	`list_ml_samples`	Shows you an overview of all the current training samples stored in your account's database.
10	`semantic_search`	Searches through your data gallery to find other samples that are conceptually similar to a provided input.

See It in Action

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

U Classify this text: 'The delivery was very late and the food was cold' using function ID 'func_123'.

Invoking ML function func_123... The prediction for your text is 'Negative Sentiment' with a confidence score of 98.45%. Shall I check if there are similar historical samples in your database?

U Search my product gallery for an image similar to 'https://example.com/shoe.jpg' using function 'func_search_99'.

Executing semantic search... I've found 3 semantically similar samples in your gallery. The top match is 'Running Shoe - Blue' (Confidence: 95.2%). Would you like the metadata for the matching samples?

U List all the machine learning functions in my Nyckel account.

Retrieving ML functions... You have 4 active functions: 'Sentiment Classifier' (func_123), 'Product Search' (func_search_99), 'Logo Detector', and 'Spam Filter'. Which function would you like to inspect?

Frequently Asked Questions

01 **How do I start classifying data using the Nyckel ML MCP?**

You must first subscribe to this MCP and provide your client credentials. Once connected, simply ask your agent to classify content by referencing a specific function ID or label.

02 **Can I find similar images without knowing keywords using Nyckel ML MCP?**

Yes, use the semantic_search tool. This feature finds samples based on conceptual similarity rather than just matching text strings, making it perfect for product discovery or image recognition.

03 **What is the difference between list_ml_labels and get_account_info?**

list_ml_labels specifically gives you all available categories your models understand. get_account_info provides broader organizational metadata about your workspace and profile.

04 **If I want to test a brand new ML model, what tool should I use in Nyckel ML MCP?**

You should start by using list_ml_functions. This shows you all currently available models, helping you decide which function to invoke_ml_function for testing.

05 **Does the Nyckel ML MCP help with data quality control?**

Absolutely. You can monitor and improve your data by listing_ml_samples and using annotate_ml_sample to manually correct or add labels to existing training records.

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 `"nyckel-ml": { "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

Nyckel ML 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 Nyckel ML. 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	Nyckel ML MCP
Server ID	`019d75e1-849f-70f2-ae45-2c6899556df8`
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/nyckel-ml.