MCP SERVER | NO CODE | CLOUD HOSTED

SingleStore MCP

Query live data and manage your entire database infrastructure.

SingleStore MCP gives your AI agent direct, read-and-write access to your SingleStore data infrastructure. Run raw SQL queries, execute semantic vector searches, list all workspaces, and audit billing usage—all from your preferred chat interface.

A+ Quality Score 100/100

sql | semantic-search | real-time-analytics | database-administration | vector-embeddings | data-infrastructure

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

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.

SingleStore MCP

6 tools available
Cloud-hosted on Vinkius

This MCP lets your AI client act like a full database administrator for your SingleStore setup. You stop bouncing between external dashboards just to check schema details or run complex search joins. Instead, you talk to your agent, which then executes the necessary commands against your live data. It handles everything from running raw SQL queries on demand to performing advanced vector similarity searches using `vector_search`. Plus, it keeps an eye on costs, letting you audit billing usage with a simple request for metrics. When integrated via Vinkius, this single connection gives your agent total control over the entire SingleStore environment, allowing deep data analysis without ever leaving your workflow.

Core Capabilities

01 — Run custom SQL queries

The agent executes raw SELECT statements against a specified database to retrieve precise data points.

02 — Perform similarity searches

It runs semantic vector searches, finding the closest matches within your dataset based on mathematical proximity.

03 — Manage and list workspaces

The agent can list all existing SingleStore workspaces and organizations associated with your account.

04 — Identify available databases

It lists all specific databases located within a given workspace ID so you know where to run queries.

05 — Audit usage and billing

You retrieve real-time metrics on your account's resource consumption and associated costs using `get_billing_usage`.

One Click on Vinkius — From Prompt to Execution

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

01 Tell your AI client what data you need (e.g., 'Show me all users who logged in last month').

02 The MCP analyzes the request, identifies the necessary tools, and executes the required actions, such as running `execute_sql` or calling `vector_search`.

03 Your agent receives structured results—the raw data tables, lists of workspaces, or usage metrics—and presents them directly to you.

The bottom line is that your AI client handles the entire round trip: understanding the need, running the query against the live database, and presenting the clean output.

Built For

Data Engineers who hate manually scripting connectivity checks; ML Scientists needing immediate access to raw vectors for model training; or Database Administrators tired of switching between monitoring dashboards.

Data Engineer

Uses the MCP to programmatically list available databases and workspaces, ensuring their ETL pipelines always point to the correct endpoints.

Machine Learning Scientist

Feeds unstructured text into the system and uses `vector_search` to find related records in the database for model fine-tuning or research.

Database Administrator (DBA)

Runs administrative checks by using `list_workspaces` and auditing costs via `get_billing_usage` without logging into a separate console.

What Changes When You Connect

01 Stop switching tabs. You run complex queries using `execute_sql` directly through conversation, getting immediate results instead of waiting for GUI forms to load.

02 Deep dive into unstructured data. Instead of manual keyword searches, use `vector_search` to find meaningful connections between records based on semantic similarity.

03 Know your limits before they bite you. Use the dedicated billing tools like `get_billing_usage` to audit costs and usage patterns instantly.

04 Manage infrastructure from one place. The agent handles listing everything—from all workspaces via `list_workspaces` to specific databases with `list_databases`—without needing admin console access.

05 Accelerate development cycles. You can list organizations or run simple checks like `list_users`, providing context to your AI client before writing complex code.

Real-World Applications

Finding the right data source for a new project

A Data Engineer needs to know if their team has already set up an environment. They ask their agent, which immediately runs `list_workspaces` and presents a clean list of available IDs, saving them manual browsing.

Investigating a sudden spike in cloud costs

A DBA suspects resource overage. Instead of pulling up the billing dashboard, they ask their agent to run `get_billing_usage`, getting an immediate, actionable metric report back in seconds.

Connecting text search to structured data

An ML Scientist has a block of unstructured user feedback. They feed it into the agent and use `vector_search` to pull up the 10 most semantically related customer records from the database, instantly grounding their research.

Verifying data availability before coding

A developer needs to know if a specific module exists. They ask the agent to `list_databases` within a known workspace ID, verifying that the schema is ready for development without running any code.

Patterns to Avoid

Manually scripting resource checks

✗ AVOID

A user writes Python code that must connect to three separate APIs—one for workspaces, one for billing, and a third for the database schema just to get basic context.

✓ INSTEAD

Let your AI client use this MCP. It coordinates calling `list_workspaces`, then running `get_billing_usage`, all in response to a single natural language prompt.

Running complex queries without scope

✗ AVOID

Writing an SQL query that references tables or databases that don't actually exist because the user forgot which workspace they were working in.

✓ INSTEAD

First, ask your agent to run `list_workspaces` and confirm the correct environment. Then execute the query using `execute_sql`.

Treating vector search like a keyword lookup

✗ AVOID

Running a simple text search that only matches exact words, missing out on contextually similar results because it ignores semantic relationships.

✓ INSTEAD

Use the `vector_search` tool. It understands meaning and finds records based on similarity vectors, giving you much richer insights than basic keyword matching.

The Right Fit

Use this MCP if your primary need is querying or auditing data already within SingleStore's ecosystem. You need the agent to run raw commands like `execute_sql` or perform specialized searches like `vector_search`. Don't use it if you are trying to build a whole new database from scratch; for that, you need an ETL pipeline tool. Also, don't rely on this MCP for user interface actions (like clicking 'Save' or 'Submit'); its job is purely data retrieval and administration. If your goal is only basic reporting without complex joins, standard BI tools might suffice, but if the query logic needs to be dynamic and based on current system state, use this MCP.

The headache of context switching in database management

Right now, checking your data environment is a manual pain. You open Dashboard A to see what workspaces exist. Then you jump to the Admin Console B to check billing limits. If you want to run a query, you have to go to Editor C, making sure you selected the right database and that your credentials are correct for that specific task.

With this MCP, all those separate hops disappear. You talk to your agent once, telling it exactly what context you need—say, 'Show me the user count and the current monthly cost.' The agent handles querying `list_workspaces`, running an `execute_sql` query, and calling `get_billing_usage` automatically. You just get the answer.

SingleStore MCP: Direct Data Access

You no longer have to copy IDs from one tab, paste them into a second script, and then manually switch to a third dashboard just to validate the scope. The agent manages these dependencies behind the scenes.

What's different is control. You don't just get *an* answer; you get an executed result derived directly from your live SingleStore infrastructure, giving you full read-and-write sovereignty over your data.

SingleStore MCP: 6 Tools for Database Control

These tools give your agent the power to interact with every part of your SingleStore environment, from running queries to checking resource usage.

#	TOOL	DESCRIPTION
01	`execute_sql`	Sends any standard SQL query to the database for execution, retrieving the resulting dataset.
02	`get_billing_usage`	Retrieves a summary report showing current resource usage and billing metrics for the account.
03	`list_databases`	Generates a list of all individual databases that exist within a specified workspace container.
04	`list_organizations`	Provides a comprehensive catalog of all organizations currently linked to the user's account.
05	`list_workspaces`	Lists every available workspace configured under the SingleStore infrastructure.
06	`vector_search`	Executes a similarity search to find records whose data vectors are closest to a provided query vector.

See It in Action

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

U List all my available workspaces.

I've successfully identified 2 distinct workspaces assigned to your account:
- **Workspace A** (ID: ws_ab12)
- **Workspace B** (ID: ws_ef56)

U List all databases within workspace ID 1234, and then find the first 5 records in 'users_db'.

The databases configured for workspace 1234 includes 'users_db' and 'logs_db'. After executing a `SELECT * FROM users LIMIT 5`, here are the detailed user rows you requested.

Frequently Asked Questions

01 **How do I use the singlestore MCP to check my monthly bill?**

You ask the agent to audit billing usage. It calls `get_billing_usage` and returns a report on your current resource consumption and associated costs.

02 **Is singlestore MCP for semantic search or just SQL queries?**

It handles both. While you can run raw SQL with `execute_sql`, the dedicated `vector_search` tool lets you perform advanced, meaning-based similarity searches against your data.

03 **What if I need to know which databases are available?**

You first use `list_workspaces` to find the container ID. Then, tell the agent to run `list_databases` using that specific workspace ID to see all contained schemas.

04 **Can I list organizations with singlestore MCP?**

Yes, simply request a list of associated accounts. The agent uses `list_organizations` to pull up a catalog of all linked organizational entities for your account.

05 **Do I need to run raw SQL queries every time with singlestore MCP?**

No. While `execute_sql` is powerful, you can also use the agent for administrative tasks like checking resource usage via `get_billing_usage`, which doesn't require running a query.

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 `"singlestore": { "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

SingleStore 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 SingleStore. 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	SingleStore MCP
Server ID	`019d7608-6b99-7129-8c92-8d643dc7228c`
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/singlestore.