

MCP SERVER

NO CODE

CLOUD HOSTED

Five9 QM MCP

Automate Quality Scoring and Review.

Five9 QM connects Quality Management workflows directly into your agent chat. It lets you pull reports, review recorded calls, and submit performance evaluations without leaving your main workspace. Manage quality assurance tasks—from checking calibration scores to generating temporary audio links—all through natural conversation.

A+ Quality Score 98.33/100

quality-assurance

performance-scoring

agent-evaluation

interaction-review

coaching

compliance-monitoring



The infrastructure that powers AI agents in the real world.



Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

Your AI Connections Run Through Vinkius Cloud

The world's largest
managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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.

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.

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.

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

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

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.

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.

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.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot 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.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

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.

Five9 QM MCP

12 tools available

Cloud-hosted on Vinkius

Managing quality assurance for a contact center is usually a headache of dashboards, spreadsheets, and clicking between systems. This MCP connects Five9 Quality Management (QM) to any agent client, letting you automate those tedious workflows using simple chat commands. Instead of logging into the QM portal just to pull metrics, your agent can list all completed evaluations or fetch detailed scores instantly. You can review interactions by retrieving full metadata for voice and digital conversations, even generating secure links to play back recordings right inside your conversation thread. If you're trying to track organizational quality trends across multiple teams, this MCP helps manage those high-level reports and keeps agents in the loop on who needs evaluating or what forms are active. By connecting Five9 QM through Vinkius, you bring powerful performance scoring tools directly into your natural workflow.

Core Capabilities

01 — Review Interactions

Fetch detailed metadata about recorded voice and digital customer interactions.

03 — Generate Recordings Links

Create temporary, secure links that allow you to play back specific audio recordings for quality review.

05 — Generate Performance Reports

Pull summary metrics that show organization-wide quality trends from QM.

02 — Manage Evaluations

Submit new performance evaluations or retrieve the full details for past assessments instantly.

04 — Analyze Forms and Calibration

List available evaluation forms or track the status of your team's calibration sessions.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/five9-qm — connect your AI agent in three steps.

- 01** Subscribe to this MCP through the Vinkius Marketplace and enter your Five9 VCC Username and Password, ensuring you have QM permissions enabled.
- 02** Send a command via your agent client (like Claude or Cursor) asking for specific quality data, such as 'List all evaluations' or 'Get metadata for interaction X'.
- 03** Your agent receives the requested data—whether it's a list of agents to evaluate or a temporary link for an audio recording—and presents it directly in your chat.

The bottom line is, you use natural language commands instead of navigating complex portals to manage all quality assurance tasks.

Built For

This connector is essential for contact center managers and QA analysts who are tired of switching between the QM platform and their daily chat workspace just to pull a score or check an agent's status. It lets them keep all performance data in one place.

Quality Assurance Analyst

Uses this MCP to quickly fetch interaction details, review recordings via secure links, and submit evaluations without leaving their primary desktop.

Contact Center Manager

Runs reports and tracks calibration progress by asking the agent for summary metrics or listing active forms, giving them a real-time view of quality scores.

Team Lead

Reviews specific recorded interactions by fetching metadata or generating recording links to coach agents efficiently during their daily workflow.

What Changes When You Connect

-
- 01 Review recordings instantly. Instead of logging into a separate player, you can request a secure link using `get_recording_link` and analyze the audio directly in your chat window for immediate coaching feedback.

 - 02 Track performance at scale. Use `list_qm_agents` to know exactly who needs an evaluation and then use `list_evaluations` to quickly pull recent scores, giving managers an instant overview of team health.

 - 03 Understand compliance gaps fast. By using `get_qm_interaction`, you retrieve all necessary metadata—call duration, topics covered, handling codes—so agents can understand *why* a score was given.

 - 04 Build consistency across teams. You don't have to manually check who has trained; simply calling `list_calibrations` shows if the team is up-to-date on scoring standards.

 - 05 Start evaluations without friction. With `create_evaluation`, you submit performance scores and feedback right from your chat interface, completing the workflow immediately.
-

Real-World Applications

Identifying a Compliance Risk

A team lead needs to check if agents handled calls about a new product line correctly. Instead of manually searching through call records, they ask their agent to `list_qm_interactions` and then retrieve the full metadata for specific interactions to confirm compliance.

Performing a Quick Score Review

A QA analyst finishes a review session and needs to record it. They use `get_evaluation` to pull all the necessary details from a past interaction, then they call `create_evaluation` to submit their final score and feedback instantly.

Onboarding New Supervisors

A manager is training new supervisors on scoring consistency. They ask the agent to `list_evaluation_forms` and `list_calibrations` to ensure the trainee knows which forms are active and when the last team calibration was held.

Auditing System Users

The IT admin needs to verify who has access rights in the QM system. They use `list_qm_users`, providing a simple roster of all system users without needing to log into the Five9 backend portal.

Patterns to Avoid

Over-relying on dashboards

X AVOID

The manager spends 20 minutes refreshing the main QM dashboard, trying to find the score for Agent Smith's last call. The data is there, but it's buried under several filters and tabs.

✓ INSTEAD

Just ask your agent client: 'Get evaluation details for Agent Smith from yesterday.' You use `get_evaluation`, and the information appears immediately without any dashboard navigation.

Copying metadata manually

X AVOID

The team lead reviews a call and then has to open three different tabs—one for the recording, one for the form, and one for the agent roster—just to compile notes.

✓ INSTEAD

First, use `get_qm_interaction` to pull all necessary interaction metadata. Then, you can `list_evaluation_forms` to reference the correct rubric structure while staying in chat.

Forgetting permissions

X AVOID

The user tries to run a report but gets an 'Access Denied' error because they forgot to enable QM permissions on their account.

✓ INSTEAD

Make sure you subscribe using your credentials that have full QM permissions enabled. If the command fails, check if you need to `list_qm_users` first to verify system access.

The Right Fit

Use this MCP if your primary bottleneck is moving between separate tools—specifically needing to get performance data (like scores or metadata) and then act on it (like submitting an evaluation)—without leaving the chat window. This is ideal for QA Analysts and Team Leads who live in their agent client.

Don't use it if you need advanced, deep-dive reporting that requires complex visual charting or cross-system data joins beyond what summary metrics provide. For pure, massive dataset analysis (like running a SQL query on 10 years of records), you might need a dedicated BI tool connector instead. However, for day-to-day operational quality checks—checking who needs scoring using `list_qm_agents`, or retrieving links with `get_recording_link`—this is exactly what you need.

The Daily Chore of Quality Review

Today, reviewing agent performance means juggling at least three different systems. You have to log into the QM portal to check scores; then switch to a separate call recording player to listen to the interaction; and finally, open another tab or document to fill out your own notes and submit the score. This constant clicking—moving from dashboard to metadata view to form selection—is where hours of productivity go.

With this MCP, all that friction disappears. You simply tell your agent what you need. For instance, if an interaction is flagged, you ask for its details using `get_qm_interaction`. Your agent pulls the full record and presents it instantly in chat. It's not about having more data; it's about getting to the answer faster.

Getting Scores and Recordings via Five9 QM

You don't have to manually search through call logs, open individual records, and then copy/paste key details into a summary sheet. You can ask your agent client to list the most recent evaluations using `list_evaluations`. If you need proof of what was said, calling `get_recording_link` gives you a temporary link right where you are.

The result is a single pane of glass experience for quality control. You keep the flow going; you don't break it. Everything—from listing available forms to submitting final scores with `create_evaluation`—happens in one place.

Five9 QM: 12 Tools for Quality Assurance

Use these twelve tools to automate every part of the QA lifecycle, from listing agents needing reviews to submitting final scores.

#	TOOL	DESCRIPTION
01	<code>create_evaluation</code>	Submits a completely new agent performance evaluation record.
02	<code>get_evaluation_form</code>	Retrieves the specific structure and questions used for an existing evaluation form.
03	<code>get_evaluation</code>	Fetches all detailed information regarding a completed agent evaluation.
04	<code>get_qm_interaction</code>	Retrieves the full metadata for any recorded voice or digital interaction.
05	<code>get_qm_summary</code>	Pulls a high-level summary report of quality metrics across the entire organization.
06	<code>get_recording_link</code>	Generates a temporary, secure link that allows direct playback of an audio recording.
07	<code>list_calibrations</code>	Displays a list of all currently active or recently completed calibration sessions.
08	<code>list_evaluation_forms</code>	Shows every quality evaluation form that is currently active and available for use.
09	<code>list_evaluations</code>	Lists the most recent completed agent evaluations, including scores and feedback summaries.
10	<code>list_qm_agents</code>	Retrieves a list of agents who are eligible to receive an evaluation score.
11	<code>list_qm_interactions</code>	Displays a record of recorded interactions available for quality review and analysis.
12	<code>list_qm_users</code>	Lists all user accounts that are configured within the QM system.

See It in Action

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

U List the last 5 completed agent evaluations.



Fetching evaluations... I found the 5 most recent evaluations, including 'Support Call Review' for Agent Smith (Score: 92) and 'Sales Interaction' for Agent Jones (Score: 85). Would you like to see the details for any of these?

U Get the recording link for interaction 'int_12345'.



Generating link... I've retrieved a secure temporary link for interaction 'int_12345'. You can listen to the recording here: https://api.five9.com/qm/v1/temp/rec_abc123. Use this for your quality review.

U Show me all available evaluation forms.



Retrieving forms... You have 3 active evaluation forms: 'General Support Rubric', 'Advanced Technical Troubleshooting', and 'Sales Compliance Checklist'. Which one should I inspect for you?

Frequently Asked Questions

01 How do I get summary metrics using Five9 QM MCP?

You use the `get_qm_summary` tool. Simply ask your agent to pull the 'QM summary report,' and it will fetch high-level organizational quality trends without you having to build a custom dashboard.

02 Can I review old calls with Five9 QM MCP?

Yes, you can. Use `list_qm_interactions` first to find the specific call ID, and then use `get_qm_interaction` to retrieve all its metadata. If needed, `get_recording_link` generates a playable audio link.

03 What if I need to score an agent right now?

You can submit evaluations directly using `create_evaluation`. This lets you record your findings and assign scores without logging into the core QM system, keeping everything in chat.

04 Does Five9 QM MCP help with team training consistency?







Absolutely. You can use `list_calibrations` to track when calibration sessions were last completed, ensuring your entire QA team is aligned on the scoring rules.

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 <code>"five9-qm": { "url": "..." }</code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → 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

Five9 QM 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](https://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 Five9 QM. 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	Five9 QM MCP
Server ID	019d759a-7862-703b-8790-a62b412912f5
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/five9-qm.