MCP SERVER | NO CODE | CLOUD HOSTED

Observe.AI MCP

Analyze call quality & conversation intelligence

Observe.AI MCP connects your AI agent directly to your contact center performance data. Get instant visibility into call transcripts, quality assurance scores, and coaching notes without leaving your workspace. Analyze every interaction—from greetings to objections—and track agent improvements using natural language prompts.

A+ Quality Score 100/100

conversation-intelligence | contact-center | quality-assurance | speech-analytics | call-transcription | performance-coaching

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

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.

Observe.AI MCP

10 tools available
Cloud-hosted on Vinkius

Connect this MCP to gain deep insight into how your customer service teams perform. You don't have to open the Observe.AI portal or manually search through spreadsheets anymore. Your AI client pulls performance data directly, allowing you to ask complex questions like, 'What was the average QA score for agents who handled billing issues last week?' The system collects everything—from full conversation transcripts to automated summaries and coaching feedback logs—and presents it in plain language. By using Vinkius, your agent gets access to this entire catalog of tools, letting you query calls, chats, and emails all from one place. This means QA Analysts can quickly check evaluation scores; Managers can monitor high-level trends during daily standups; and Coaches can verify improvement history instantly.

Core Capabilities

01 — Review all service interactions

List every call, chat, or email processed by the platform, along with metadata.

02 — Retrieve full conversation text

Pull the complete text transcript for any specific interaction so you can review details instantly.

03 — Assess agent quality scores

Access formal quality assurance evaluation forms, individual scores, and performance metrics.

04 — Identify key conversation moments

List specific business moments identified by the AI, such as greetings or customer objections, across multiple interactions.

05 — View summarized call themes

Read automated summaries that distill the main topics discussed in recent conversations.

06 — Track coaching history

List and review records of agent coaching sessions and feedback given by supervisors.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP on Vinkius and input your Observe.AI API Key (Bearer Token).

02 Connect your preferred AI client, like Cursor or Claude, to the catalog.

03 Ask your agent a question—for example, 'What were the top three objections raised by customers last month?'—and it pulls the data.

The bottom line is you get instant answers about your contact center performance without logging into multiple dashboards or systems.

Built For

This MCP is for Operations Managers and Supervisors who are tired of manually aggregating data from disparate portals. If you spend time copying transcripts, hunting down QA scores, or summarizing trends for a meeting, this is for you.

Contact Center Manager

Monitoring high-level performance and running daily standups by asking the agent for AI-generated summaries of recent interactions.

QA Analyst

Quickly checking evaluation scores or reading full transcripts for specific interaction IDs without having to open the main Observe.AI portal.

Training Supervisor

Verifying agent progress and reviewing coaching session history using natural language prompts while talking to their team.

What Changes When You Connect

01 Stop hunting for transcripts. Use get_interaction_transcript to pull the full text of any chat or call in seconds, letting you review conversations without leaving your workflow.

02 Track performance trends instantly. Rather than digging through dozens of reports, list_qa_evaluations gives you a clean summary of quality scores across teams and time periods.

03 Automate daily reporting. Use list_interaction_summaries to get immediate high-level overviews of what customers are talking about—great for quick manager updates.

04 Spot training gaps fast. By calling list_coaching_sessions, supervisors can see exactly when an agent was coached and on what topics, proving progress or identifying recurring weakness.

05 Go deeper than scores. Calling list_interaction_moments lets you pinpoint *why* a call failed—did they miss the 'Closing' moment? Did they fail to acknowledge the customer's 'Objection'?

06 Manage personnel data easily. Use list_workspace_users to quickly verify who is on your team and who needs access, all without logging into the main admin portal.

Real-World Applications

Investigating a bad customer experience.

A manager hears about a repeat complaint. They ask their agent to pull the full transcript using get_interaction_transcript for all interactions involving that customer in the last week. This immediately highlights patterns and shows exactly where the service broke down.

Preparing for team reviews.

A coach needs proof of improvement. They use list_coaching_sessions to pull a history of past feedback, then ask the agent to compare that progress against recent QA scores using get_evaluation_details.

Quickly understanding market shifts.

The team needs to know if billing issues are spiking. They prompt the agent for list_interaction_summaries, which quickly reveal 'Billing Inquiries' as the dominant theme across all interactions this month.

Auditing compliance failures.

A QA analyst suspects a team is missing mandatory disclosures. They use list_interaction_moments to filter for instances where the required 'Compliance Statement' was not recorded, providing actionable data points.

Patterns to Avoid

Downloading everything into Excel.

✗ AVOID

Exporting 50 reports containing transcripts, scores, and summaries into a single spreadsheet. It takes hours of cleanup to make sense of the mixed data types and missing dates.

✓ INSTEAD

Instead, let your agent pull specific metrics directly, like calling list_qa_evaluations for all agents who scored below 85%, then using get_interaction_details to grab just their names. Focus on targeted questions instead of mass downloads.

Asking vague 'How' questions.

✗ AVOID

Prompting the agent with, 'Tell me about our performance.' The response is a massive, unfilterable blob of text that doesn't help you narrow down the actual problem area.

✓ INSTEAD

Be specific. Ask, 'What were the most common objections and their associated scores from the last 30 days?' This directs the agent to use list_interaction_moments alongside list_qa_evaluations.

Forgetting context.

✗ AVOID

Needing to know if an agent was trained on a specific process. You might only ask for 'agent history' without specifying the type of training record.

✓ INSTEAD

Always specify what you need. To check training, use list_coaching_sessions and filter by topic name or date range.

The Right Fit

Use this MCP if your primary job involves reviewing conversation details—transcripts, scores, summaries—to improve customer service processes. You need to ask questions like, 'Why did that call fail?' or 'What was the theme of these 10 chats?' This tool excels at pulling deep analytics and structural data (like list_interaction_moments). Don't use this if your main goal is simple ticketing management; for basic record keeping, a standard CRM connector works better. Also, don't expect it to generate new content; it analyzes existing performance records. If you just need to write internal documentation or draft replies, stick to pure writing AI tools instead.

The Manual Burden of Call Quality Review

Every week, the process is the same: you open the portal. You filter by date. You manually pull transcripts for agents who missed a metric. Then you copy the scores into a spreadsheet and cross-reference them with coaching forms to figure out where the training actually needs to happen. It's tedious, slow, and it takes hours just to gather enough data to make one decision.

With this MCP, your agent handles the grunt work. You simply ask: 'Show me all interactions from last week that scored under 80% for compliance.' The system retrieves the list of failing calls and even provides specific details using get_interaction_details. You immediately get a actionable, focused report.

Get Full Visibility With Observe.AI MCP

No more jumping between the transcript viewer, the QA score tab, and the coaching log. These key pieces of information are separate systems that usually require three different logins and five clicks each.

Now, your agent pulls all this context together. You ask for a summary, and it aggregates data from list_interaction_summaries, get_evaluation_details, and even lists related moments using list_interaction_moments—all in one query.

Observe.AI MCP: 10 Tools for Service Intelligence

These tools allow your agent to execute specific tasks like pulling transcripts, listing evaluations, and gathering interaction metadata, giving you granular control over data retrieval.

#	TOOL	DESCRIPTION
01	`get_evaluation_details`	Retrieves specific quality assurance evaluation details for an interaction.
02	`get_interaction_details`	Gets general metadata and information about a specific customer interaction.
03	`get_interaction_transcript`	Pulls the full text transcript of a recorded call or chat conversation.
04	`list_coaching_sessions`	Lists all documented coaching sessions for a specific agent.
05	`list_qa_evaluations`	Retrieves a list of all available quality assurance evaluations.
06	`list_evaluation_forms`	Lists the specific forms used for QA evaluation.
07	`list_interactions`	Retrieves a list of recent contact center interactions, including calls and chats.
08	`list_interaction_moments`	Lists key business moments (like 'Greeting' or 'Objection') identified by the AI within interactions.
09	`list_interaction_summaries`	Provides a list of automated, high-level summaries for recent interactions.
10	`list_workspace_users`	Retrieves a directory listing of agents and administrative users in the Observe.AI workspace.

See It in Action

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

U List all recent call interactions from today.

Retrieving interactions... I found 12 calls from today. Notable ones include 'Interaction #98765' (Duration: 5:30, Sentiment: Positive) and 'Interaction #98766' (Sentiment: Frustrated). Would you like the transcript for the frustrated call?

U What is the QA score for interaction ID 'int_12345'?

Fetching evaluation details... The interaction int_12345 received a total QA score of 95%. It scored perfectly on 'Compliance' and 'Empathy', with a minor deduction in 'Resolution Speed'. Shall I list the full evaluator feedback?

U Show me the AI summaries for our latest interactions.

Retrieving AI summaries... I found 5 recent summaries. The overall theme focuses on 'Renewals' and 'Billing Inquiries'. One summary identifies a recurring issue with 'Portal Login'. Would you like the detailed summary for that specific case?

Frequently Asked Questions

01 **How does Observe.AI MCP handle transcripts?**

You can retrieve the full text transcript for any call or chat interaction by calling get_interaction_transcript. This gives you the complete conversation history immediately in your agent's response.

02 **Can I check historical QA scores using Observe.AI MCP?**

Yes, you can list all available quality assurance evaluations using list_qa_evaluations to see a record of past scoring efforts and trends.

03 **What is the best way to analyze agent performance with Observe.AI MCP?**

Start by listing interactions using list_interactions, then ask for get_interaction_details on any specific ID. This gives you core metadata necessary to understand context before diving into scores.

04 **How do I find out what customers are complaining about?**

Ask the agent to use list_interaction_summaries or list_interaction_moments. These tools automatically identify recurring themes and key moments like 'Objection' across many calls.

05 **Does Observe.AI MCP help with coaching records?**

Yes, you can use the list_coaching_sessions tool to pull a history of agent training sessions and track when specific feedback was given.

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 `"observeai": { "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

Observe.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 Observe.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	Observe.AI MCP
Server ID	`019d75e1-cd6e-735f-b5b0-e68ce4fb3c64`
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/observeai.