

MCP SERVER

NO CODE

CLOUD HOSTED

8x8 Contact Center MCP for AI Agents

Real-time queue status and historical agent activity auditing

The 8x8 Contact Center MCP connects your AI agent directly to real-time call center analytics. It lets you audit live queue status, review historical agent performance logs, and analyze long-term contact center trends through simple conversation.

F Quality Score 3.6/100

contact-center

ccaas

call-analytics

queue-management

agent-performance

real-time-metrics



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.

8x8 Contact Center MCP

3 tools available

Cloud-hosted on Vinkius

This integration turns your AI client into a supervisor for your 8x8 Contact Center. You stop digging through complex dashboards and start asking questions in plain language. Your agent provides a direct window into call center operations, whether you need an instant pulse check on live calls or a deep dive into past employee activity.

It pulls current metrics, letting you see exactly what's happening with queues and agents right now. You can also audit historical records; for example, retrieving specific logs of agent interactions by date allows quality assurance teams to quickly pull metadata for performance reviews. Need long-term context? Analyzing queue performance over time shows the bigger picture of contact center trends. Getting this data all managed through your AI client means you stay focused on solving problems, not navigating interfaces. You connect everything via Vinkius, giving you access to thousands of other services in one place.

Core Capabilities

01 — Get real-time operational metrics

Retrieve current statistics for all active queues and agents instantly.

02 — Audit historical agent logs

Review specific records of past agent interactions, filtered by date to pinpoint resolution details.

03 — Analyze queue performance over time

Access aggregated data showing how queues have performed historically, identifying long-term bottlenecks or improvements.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/8x8-contact-center — connect your AI agent in three steps.

- 01 Subscribe to this MCP and enter your 8x8 API Key and Client Secret.
- 02 Connect the service to your preferred AI client (Claude, Cursor, etc.).
- 03 Use natural language prompts within your chat interface to audit metrics or pull performance data.

The bottom line is, you talk to your agent about your contact center needs, and it pulls the relevant 8x8 operational data for you.

Built For

This MCP targets operations leaders who are tired of clicking through multiple dashboards just to get a simple status update. It's built for those who need an immediate, conversational view into call center health, whether they manage staffing or quality assurance.

Contact Center Supervisor

Needs to perform real-time pulse checks on queue health and agent availability without navigating complex status dashboards.

Operations Manager

Uses the MCP to audit historical interaction logs and performance trends, helping them optimize staffing levels across departments.

Quality Assurance Team Lead

Quickly pulls specific agent interaction metadata using date filters for immediate performance review and compliance checks.

What Changes When You Connect

- 01 Immediate operational awareness: Instantly check live call volumes and agent availability using `get_realtime_metrics`, eliminating the need to switch tabs or load separate dashboard pages.

-
- 02 Deep audit capability: List specific agent interactions by date range via `list_agent_interactions`. This is crucial for QA teams needing metadata on past calls for compliance checks.

 - 03 Strategic planning data: Analyze long-term contact center trends by calling `list_queue_metrics`, helping managers predict staffing needs and identify consistent bottlenecks.

 - 04 Conversational oversight: Audit agent status and queue health using simple prompts instead of complex filtering menus. The AI does the heavy lifting.

 - 05 Focus on insights, not clicks: Your team gets actionable data points directly in the chat window, keeping your workflow centralized and fast.
-

Real-World Applications

Investigating a sudden spike in wait times

An Ops Manager notices slow service. They ask their agent what the current live status of the queues is. The MCP uses `get_realtime_metrics` to confirm that the 'Billing' queue has 30 calls waiting, immediately pointing to an operational bottleneck.

Evaluating seasonal staffing changes

The manager wants to know if the 'Tech Support' queue is trending poorly over quarters. They use `list_queue_metrics` to pull aggregated historical performance data, proving that wait times spike consistently in Q4.

Reviewing a specific employee's performance

A QA lead needs details on three specific calls from last month. They ask their agent to list agent interactions filtered by date. The MCP executes this and provides the metadata for those exact calls, saving hours of manual log searching.

Patterns to Avoid

Trying to manually compare logs

X AVOID

An agent copies and pastes date ranges into a separate reporting tool just to find the call duration for yesterday's interactions. This is slow, tedious, and prone to copy-paste errors.

✓ INSTEAD

Use your AI client with `list_agent_interactions`. You simply tell it: 'Show me all agent calls from last Tuesday.' It handles the date filtering and data retrieval in one step.

Ignoring real-time status checks

X AVOID

A supervisor walks over to a desk, only to find out that a queue is backed up because no one checked the live metrics. They waste time manually checking the system.

✓ INSTEAD

Use your AI client with `get_realtime_metrics` first thing in the morning. You get an instant, accurate pulse check on agent status and current call volumes.

Confusing historical data with trends

X AVOID

Looking at one day's worth of metrics (`list_agent_interactions`) and assuming it represents a long-term problem. You can't tell if the issue is isolated or systemic.

✓ INSTEAD

Always look to `list_queue_metrics` for long-term context. This function provides aggregated performance data, helping you spot sustained trends rather than one-off spikes.

The Right Fit

Use this MCP if your biggest pain point is the time spent switching between dashboards or manually querying logs to get an operational status update. You need a central conversational hub that can access live metrics (`get_realtime_metrics`), historical records (`list_agent_interactions`), and long-term trends (`list_queue_metrics`). Don't use this if you only need to write reports based on data already exported into a spreadsheet, because the MCP is designed for querying the data itself. If your goal is simply reporting or visualization using pre-existing files, look at general BI tools; but if the data lives within 8x8 and needs conversational access, this is what you want.

Managing 8x8 Contact Center Metrics with AI Agents

Right now, checking your team's operational status means jumping between the main dashboard, the agent roster, and separate reporting tools. You have to click through 'Support Queue,' then find the live call count, then check individual agents for their 'Available/Busy' status, and repeat that process every hour just to get a quick pulse check.

With this MCP, you ask your AI agent: 'What is the current status of the support queues?' It immediately runs the necessary checks and gives you the live metrics in plain text. You get an instant summary without leaving the chat interface.

Auditing Agent Activity Using 8x8 Contact Center with AI Agents

If a customer complains about service quality, you currently have to open the interaction log system, manually find that agent's ID, and then use date filters to narrow down call records. This process is slow, requires multiple logins, and wastes valuable time during an investigation.

Now, your AI agent handles it all. You simply prompt: 'List all interactions for Agent X between June 1st and June 5th.' The MCP retrieves the complete history of metadata and statuses instantly, giving you actionable data without any manual searching.

3 Tools for 8x8 Contact Center Queue & Performance Analytics

Use these tools to get real-time metrics, audit agent activity history, or analyze long-term queue performance data directly through your AI client.

#	TOOL	DESCRIPTION
01	<code>get_realtime_metrics</code>	Pulls live statistics on the current status of all active contact center queues and agents.
02	<code>list_agent_interactions</code>	Retrieves historical records of agent interactions, allowing filtering by date to review call resolution metadata.
03	<code>list_queue_metrics</code>	Accesses aggregated performance data detailing how specific queues performed over a selected time period.

See It in Action

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

U What's the current status of our queues?



Live Status Report

- **Support Queue:** 15 calls waiting. (Priority: High)
- **Sales Queue:** 2 calls waiting. (Priority: Medium)

Agent Load Breakdown:

- Support Agents: 8/10 available.
- Sales Agents: All busy. Need to reallocate resources.

U Show me interaction logs for last week's billing calls.



Interaction Logs (Last Week)

AGENT	DATE RANGE	CALL DURATION	STATUS
J. Smith	6/10 - 6/12	8m 30s	Resolved (Billing)
M. Lee	6/11 - 6/12	4m 15s	Escalated (Fraud)

Total interactions found: 45. Need a deeper dive on call duration?

U How has the general queue performed in Q1?



Quarterly Queue Performance Analysis

The 'General' queue showed a steady increase in average wait time from January to March.

- Q1 Start Avg Wait: 32 seconds (Service Level: 90%)
- Q1 End Avg Wait: 45 seconds (Service Level: 85%)

This trend suggests staffing needs increased capacity by at least 15% for Q2.

Frequently Asked Questions

01 How does the 8x8 Contact Center MCP help with daily queue status checks?

It gives you an instant, natural language summary of your live call center metrics. Instead of loading a dashboard, you ask your agent what's happening right now and get immediate details on waiting calls and agent availability.

02 Can I use the 8x8 Contact Center MCP to check past employee performance?

Yes, it lets you audit historical records. You can request a list of interactions for any given date range, retrieving crucial metadata like call duration and resolution status without manual log searching.

03 Is the 8x8 Contact Center MCP useful for predicting staffing needs?

It helps by providing long-term trends. By analyzing historical queue performance, you can spot patterns—like consistent wait time spikes every quarter—to justify better resource allocation.

04 What kind of data can I get from the 8x8 Contact Center MCP?







You get live metrics for queues and agents, historical lists of agent interactions with metadata, and aggregated performance data showing how queues trend over time. It covers everything you need to audit operations.

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>"8x8-contact-center": { "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

8x8 Contact Center 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 8x8 Contact Center. 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	8×8 Contact Center MCP
Server ID	019d7545-44f1-72a4-b477-f3a488620225
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/8x8-contact-center.